P2V Conversion Tools p. 59

Windows® IT Pro

A PENTON PUBLICATION

JANUARY 2011 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

SBS moves to the Cloud p. 21

Interview with Microsoft SBS exec
Kevin Kean p. 15

Troubleshooting Active Directory p. 26

Master PowerShell Functions p. 29

Desktop Virtualization and VDI Architecture p. 33

Monitor Key Performance Indicators with System Center Operations Manager p. 39

Secure Client Access Servers p. 44

Maintenance for SharePoint Databases p. 49

New Column!

Sean Deuby on Enterprise Identity p. 15

Smarter technology for a Smarter Planet:

What 99.9% system uptime means to a kilo of gold.

It means that the futures contract for that gold can trade instantly and more securely. The Dubai Gold & Commodities Exchange (DGCX) has maintained their complex network of worldwide members for four years without a single security breach due to malware, and without any unplanned downtime. The DGCX worked with IBM Security Solutions to help implement an intrusion prevention system that builds security into every aspect of their online trading services and proactively adapts to ever-evolving threats. A smarter business is built on smarter software, systems and services.

Let's build a smarter planet. ibm.com/exchange

IBM

A data visualization of the settlement prices for gold, silver and other commodities from March 1 to September 1, 2010.

C O N T E N T S

Windows IT Pro
A PENTON PUBLICATION
JANUARY 2011
VOLUME_17
NO_1

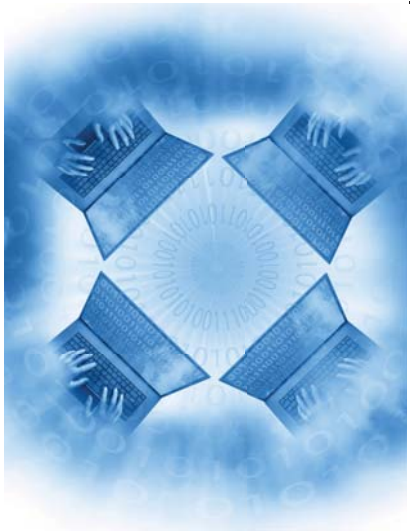

COVER STORY

FEATURES

INTERACT

IN EVERY ISSUE

Access articles online at www.windowsitpro.com. Enter the article ID (located at the end of each article) in the InstantDoc ID text box on the home page.

COLUMNS

C O N T E N T S

PRODUCTS

Windows IT Pro

EDITORIAL

Editorial and Custom Strategy Director
Michele Crockett mcrockett@windowsitpro.com

Editor in Chief
Amy Eisenberg amy@windowsitpro.com

Senior Technical Director
Michael Otey motey@windowsitpro.com

Technical Director
Sean Deuby sdeuby@windowsitpro.com

Senior Technical Analyst
Paul Thurrott news@windowsitpro.com

Industry News Analyst
Jeff James jjames@windowsitpro.com

Custom Group Editorial Director
Dave Bernard dbernard@windowsitpro.com

Developer Content
Anne Grubb agrubb@windowsitpro.com

Exchange & Outlook
Brian Winstead bwinstead@windowsitpro.com

Networking, Storage, Hardware
Jason Bovberg jbovberg@windowsitpro.com

SharePoint
Caroline Marwitz cmarwitz@windowsitpro.com

SQL Server
Megan Keller mkeller@windowsitpro.com

Systems Management, Virtualization, Windows OS
Zac Wiggy zwiggy@windowsitpro.com

Editorial Web Architect
Brian Reinholz breinholz@windowsitpro.com

CONTRIBUTORS

SharePoint and Office Community Editor
Dan Holme danh@intelliem.com

Senior Contributing Editors
David Chernicoff david@windowsitpro.com
Mark Joseph Edwards mje@windowsitpro.com
Kathy Ivens kivens@windowsitpro.com
Mark Minasi mark@minasi.com
Paul Robichaux paul@robichaux.net
Mark Russinovich mark@sysinternals.com

Contributing Editors
Alex K. Angelopoulos aka@mvps.org
Sean Deuby sdeuby@windowsitpro.com
Michael Dragone mike@mikerochip.com
Jeff Fellinge jeff@blackstatic.com
Brett Hill brett@iisanswers.com
Darren Mar-Elia dmarelia@windowsitpro.com
Tony Redmond 12knocksinna@gmail.com
Ed Roth eroth@windowsitpro.com
Eric B. Rux ericbrux@whshelp.com
John Savill john@savilltech.com
William Sheldon bsheldon@interknowlogy.com
Randy Franklin Smith rsmith@montereytechgroup.com
Curt Spanburgh cspanburgh@scg.net
Orin Thomas orin@windowsitpro.com
Douglas Toombs help@toombs.us
Ethan Wilansky ewilansky@windowsitpro.com

ART & PRODUCTION

Production Director
Linda Kirchgesler linda@windowsitpro.com

Senior Graphic Designer
Matt Wiebe matt.wiebe@penton.com

ADVERTISING SALES

Publisher
Peg Miller pmiller@windowsitpro.com

Director, International and Agency Services
Don Knox don.knox@penton.com

Business Development Director
Kerry Gates kerry.gates@penton.com

EMEA Managing Director
Irene Clapham irene.clapham@penton.com

Director of IT Strategy and Partner Alliances
Birdie J. Ghiglione birdie.ghiglione@penton.com 619-442-4064

Online Sales and Marketing Manager
Dina Baird Dina.Baird@penton.com

Key Account Director
Chrissy Ferraro christina.ferraro@penton.com 970-203-2883

Account Executives
Barbara Ritter barbara.ritter@penton.com 858-367-8058
Cass Schulz cassandra.schulz@penton.com 858-357-7649

Client Project Managers
Michelle Andrews 970-613-4964
Kim Eck 970-203-2953

Ad Production Supervisor
Glenda Vaught glenda.vaught@penton.com

MARKETING & CIRCULATION

Customer Service service@windowsitpro.com

IT Group Audience Development Director
Marie Evans marie.evans@penton.com

Marketing Director
Sandy Lang sandy.lang@penton.com

CORPORATE

PENTON

Chief Executive Officer
Sharon Rowlands Sharon.Rowlands@penton.com

Chief Financial Officer/Executive Vice President
Nicola Allais Nicola.Allais@penton.com

TECHNOLOGY GROUP

Senior Vice President, Technology Media Group
Kim Paulsen kpaulsen@windowsitpro.com

WRITING FOR *WINDOWS IT PRO*
Submit queries about topics of importance to Windows managers and systems administrators to articles@windowsitpro.com.

LIST RENTALS
Contact MeritDirect, 333 Westchester Avenue, White Plains, NY or www.meritdirect.com/penton.

REPRINTS
Diane Madzelonka, Diane.madzelonka@penton.com, 216-931-9268, 888-858-8851

Those who can envision a plausible future that's brighter than today will earn the opportunity to lead.

Ray Ozzie October 28, 2010

The world has changed.

- You need to know about your SQL Servers, whenever, wherever.
- You're available 24/7, but your home life matters. It matters a lot.
- You're always carrying an internet-connected device.

18 months ago, Red Gate started the biggest project in its history to create the future of SQL Server monitoring.

Find out whether we've earned the opportunity to lead:
www.thefutureofmonitoring.com

sql MONITOR™

redgate®
ingeniously simple tools

Crockett

"Don't make your boss drag you kicking and screaming into new technology. Ideally, you'll be the one setting meetings to brief your company leaders on new technology that can help the business run smoothly and cost-effectively."

5 Ways to Master Your 2011 IT Budget

Make the business better with smart IT spending

Got your IT budget for 2011? Or is your budget still in process as your company evaluates its business objectives? Although the economy is slowly improving, many IT organizations are grappling with developing budgets that eke the best value out of every last dollar. IT pros need to mitigate system degradations that likely occurred during the last two years of budget pullbacks, ensure that business processes run smoothly, anticipate company strategy shifts that will require different systems, and keep costs down. All in a day's work, right?

Building a budget that's a win for company leaders and the IT organization starts with a firm understanding of business drivers and the technology. Jeff James discusses the business side of building the technology budget in his Business Technology Perspectives column. To get some input on the IT budget process, we talked with a few members of the team at eTek Global, a consulting company based in Overland Park, Kansas, that works with small and medium-sized businesses. Ameet Phadnis, a partner and principal consultant for eTek Global, noted that most of their clients' technology investments in 2011 will focus on business process automation and business intelligence systems. Cindi Reding, vice president of business development, added that implementing SharePoint is a primary focus for many companies because it can put business data in the hands of the information workers, freeing up the IT staff for other projects. Phadnis and Reding, along with Leonard Mwangi, also a partner and consultant, offered useful observations about how IT pros can drive a successful budget process.

Understand the business. The need to understand your company's business is a theme that just won't die, but it's a concept that many IT pros still haven't mastered, according to Reding. "You have to know the business, you have to understand the business processes, and you have to be able to talk to the business leaders on a business level and not a technical level." Reding says that until IT teams include business-oriented technologists, "we're never going to get past this mindset where IT is seen just as a cost center. You need to have a group of people who can actually show a benefit rather than just a bunch of tools that are hard to implement."

Evaluate the latest technology—before your boss does. IT pros everywhere are familiar with the scenario in which the business leader hears about a new, innovative technology on the radio while driving to work and calls a meeting that afternoon to rally the IT team to buy and implement the game-changing product. To add to the fun, the new technology might be on a completely different platform from existing technology. In this case, your best recourse is to focus on the business problem that the company leader is trying to accomplish rather than the specific product or technology. If you can't offer an immediate critique of the technology and propose an action plan, then acknowledge the business leader's intent to improve business processes. Commit to researching a solution and following up with an action plan. Don't make your boss drag you kicking and screaming into new technology. Ideally, you'll be the one setting meetings to brief your company leaders on new technology that can help the business run smoothly and cost-effectively.

Take charge of the proof of concept. When the business needs new systems or tools, your engagement from the outset as an IT professional can help determine the success of the project. Seize the reins of the proof-of-concept phase so that technology considerations are built into the project planning and feasibility assessments.

Know your licensing needs. The most common pitfall in the nuts-and-bolts of budget building is misunderstanding the licensing requirements for products. "Licensing almost always costs more than you think" once you've done a true calculation of costs, Reding said. Phadnis added that hardware-related costs also have the tendency to expand. "Everyone talks about creating virtual machines as a way to cut costs, but they don't always consider whether the physical machine can handle the load," Phadnis said.

Insist on a training budget. One reason that IT pros are dragged kicking and screaming into new technology is because they lack the skills required by the new platform. When your company is investing in new systems and tools, insist that the budget include training for you and your staff. Well-executed boot camps can quickly pay for themselves by reducing false starts in new technology implementation.

How does your IT budget look for 2011? Do you have new projects on the horizon? Send your thoughts to me via email at michele.crockett@penton.com. And follow me on Twitter @michelecrockett.

InstantDoc ID 129062

MICHELE CROCKETT (michele.crockett@penton.com) is editorial strategy director of Penton Media's IT and developer publications, including *DevProConnections, Windows IT Pro, SharePoint Pro Connections, SQL Server Magazine,* and *Connected Planet*.

James

"When business leaders start making IT decisions without the input of IT staff, it really can be a recipe for disaster."

IT Budgeting Tips

Cooperate with IT for business success

The IT budgeting process is an annual tradition for most businesses and helps IT staff and business owners come together to plot out their IT spend over the next 12 months. To get some perspective (and some advice) on this process, my colleague Michele Crockett and I spoke with some IT budgeting experts at eTek Global. Michele Crockett discusses budgeting on the IT side of the house in her IT Pro Perspectives column. I'll focus more on the business side of the discussion in this month's Business Technology Perspectives.

Cindi Reding is the former CIO of Penton Media and is currently the vice president of business development at eTek. Reding mentions that IT departments and business leaders would be well-served by working closely together during the budgeting process, mainly to ensure that both IT and business needs are being met. "When business leaders start making IT decisions without the input of IT staff, it really can be a recipe for disaster," says Reding. "Both IT and management need to avoid having IT go down one track and business stakeholders going down another."

Reding stresses that the right working conditions really need to be set at the top of the company, with senior leadership committed to making sure that business leaders and IT work together. Getting the executive team involved and active in pushing for cooperation between IT and other groups can reap financial benefits by avoiding costly mistakes, particularly when a new IT platform or service is deployed that doesn't meet the needs of the business, or doesn't integrate well with existing IT systems.

"Another thing that business leaders can do to help the process along is to not get distracted by bright and shiny objects," Reding says. "I've heard of several cases where a senior executive read something about a new IT technology or product while reading a newspaper or magazine during a plane flight, then encouraged the IT department to adopt the technology without knowing enough about the technical ramifications of adopting it."

On the IT side of the house, Reding says that IT pros would be well-served by fully understanding the business needs of the company during the budgeting process, and to focus on providing solutions that are truly a benefit for the business. "IT departments have to demonstrate the benefits those new IT purchases will bring to the company," Reding says. "IT departments really need to make sure that they're not being seen as just a cost center to the rest of the company."

Reding has some additional suggestions for ensuring a smooth IT budgeting process:

Invest in training. Business leaders should be open to additional funds for training IT staff in the use of the latest products and technologies. Technologies like virtualization and cloud computing can have positive financial results for the companies that adopt them, but only if the IT staff has been effectively trained in how to deploy, manage, and secure them.

Conduct proof-of-concept and deployment tests. For more complex product and platform deployments, work with vendors to create proof of concepts that show how the adoption of new technology may work with an existing IT infrastructure. This approach can help avoid costly deployment problems and help get the most out of limited IT budgets. Many IT professionals have tales to tell about costly deployments of new tools or services that were reversed due to incompatibilities with legacy systems or hidden requirements, so proof of concepts and test deployments can help you avoid costly (and potentially career-ending) mistakes.

Many IT professionals have tales about costly deployments of new tools or services that were reversed due to incompatibilities with legacy systems or hidden requirements.

Do you have any suggestions for how IT and business leaders can work together during the IT budgeting process? Pass along your suggestions to jeff.james@penton.com or send a Tweet to my attention on Twitter @jeffjames3.

InstantDoc ID 129064

JEFF JAMES (jeff.james@penton.com) is industry news analyst for *Windows IT Pro*. He was previously editor in chief of Microsoft *TechNet magazine*, was an editorial director at the LEGO Company, and has more than 15 years of experience as a technology writer and journalist.

The Conversation Begins Here

mobile connections

April 17-20, 2011
Las Vegas, Nevada

iPhone · Windows Phone 7

Questions Answered • Strategy Defined • Relationships Built


Jay Freeman
CYDIA


Aaron Hillegass
BIG NERD RANCH


Ilja Laurs
GETJAR


Joe Belfiore
MICROSOFT

Top Reasons to Attend

- Acquire new skills across platforms—iOS/iPhone • Android • Windows Phone 7 • BlackBerry
- Learn how to leverage cloud-based services for mobile app delivery
- Learn how to market and sell your apps
- Get the latest market predictions
- Compare security solutions for mobile app delivery
- **Network with your peers, IT professionals, carriers and a wide range of mobile infrastructure, product and service vendors!**

EARLY BIRD DISCOUNT

Register by **February 1** *and book a minimum of three nights at Bellagio and you'll receive a $100 Bellagio Gift Certificate!*

PLUS! *you'll get a special attendee rate at Bellagio of $149.*

BOOT CAMP

Attend one of these pre-conference workshops:

- iOS/iPhone
- Android
- Windows Phone 7
- BlackBerry

Our instructors will walk you through the fundamentals of native application programming, from mobile application architecture to deployment to the online market.

CO-LOCATED WITH:

cloud connections

virtualization connections

Register for one event and gain access to sessions at all three events at **no additional cost**.

REGISTER: *MobileConnectionsEvent.com or 800.505.1201*

Thurrott

"The more I use Office 365, the more I become convinced that this is the future of Microsoft all tied up in one neat little product bundle."

What You Need to Know About Microsoft Lync 2010, Office 365, IE 9, and Windows Phone Carriers

It's 2011 already? And I was just getting used to saying twenty-ten. Here's what you need to know about Microsoft Lync 2010, an enterprise service whose time has come; why Office 365 makes sense; what Internet Explorer (IE) 9 offers that no other browser will ever offer on Windows; a broken promise with Windows Phone 7, and why Kinect for Xbox 360 will change how humans interact with computers.

Lync's Time Has Come

Microsoft Lync 2010 Standard and Enterprise will be available by the time you read this, and Lync Online is coming soon as part of Office 365 (see below). This hybrid approach to mainstream Microsoft servers is going to be quite common going forward, with the software giant offering both on-premises and hosted versions of its servers, giving customers a choice of where to put their infrastructure.

More specifically, Lync is the next-generation version of what used to be called Office Communications Server (OCS). It provides enterprise-class presence, instant messaging (IM), audio- and video-conferencing, and more, in a package that naturally integrates with other Microsoft products, especially Exchange, SharePoint, and Outlook. There's a client application that's part of the Office suites and new PBX-like capabilities that will let us finally step out from the shadow of the ancient phone systems on which many businesses still rely.

OCS was, perhaps, a bit ahead of its time. But I recommend evaluating Microsoft Lync.

Office 365

Last month, I wrote about Office 365, and about my belief that this new offering is all about putting Microsoft's most popular productivity servers and applications on a subscription payment plan. Now that I've been using a beta version of the service, I feel that this is still the case. But that doesn't mean Office 365 doesn't make sense for Microsoft's customers as well. In fact, it makes plenty of sense.

Office 365 is being made available in a wide range of product versions, but Microsoft neatly divides them into two main categories, Office 365 for Small Businesses and Office 365 for Enterprises. The small business version is aimed at businesses of one to 50, though most businesses will want to consider the upper-level enterprise-oriented options after they exceed 25 employees or so.

For those small businesses, Office 365 is pretty compelling. I'm talking some nice features here:

- hosted Exchange 2010 with 25GB of storage space for each mailbox
- self-service team collaboration websites via hosted SharePoint 2010
- IM, presence, and online meetings with audio- and video-conferencing through Lync Online 2010
- private versions of the Office Web Apps (web versions of Word, Excel, PowerPoint, and OneNote) as well as Outlook Web App for email, calendar, tasks, and contacts management (and if employees have desktop versions of Office 2010, they're free to use those applications as well)
- simple, central web portal, with no IT required
- 99.9 percent uptime SLA

The cost for this service? $6 per user per year. Think about that for a second. Google's small business offering is a bit less expensive, about $4.20 per month, though you must pay for a year at a time. But Office 365 provides you with real Exchange, real SharePoint, and Lync-based conferencing capabilities that Google can't touch. There's no comparison.

Office 365 for Enterprises is actually several different product versions, and businesses are free to mix and match, providing different employees with different levels of service, functionality, and, of course, pricing. The basics are the same as the small business offering, but enterprise customers also get 24 x 7 phone support, single sign-on (SSO—and, optionally, federation) with on-premises Active Directory (AD), and the current (2010) version of Office Professional Plus, which includes desktop-based versions of Word, Excel, PowerPoint, Outlook with Business Contact Manager, OneNote, Publisher, Access, InfoPath, SharePoint Workspace, and Lync (client).

This version costs $24 per month per user, but there are many other enterprise offerings, including a kiosk offering (for light email and SharePoint usage). The prices vary accordingly, with some coming in even below the Small Business version: The kiosk offering is just $2 per user per month.

The more I use Office 365, the more I become convinced that this is the future of Microsoft all tied up in one neat little product bundle. The only thing missing is an Application Virtualization (App-V)-style remote application deployment model for the local Office applications, though one has to wonder if that isn't in the

plans for version two. But even in its current form, Office 365 is proof positive that Microsoft's plan to move to the cloud isn't just viable—it's a good one.

Internet Explorer 9 Performance

When Microsoft introduced IE 9 at PDC 2009, it promised to better adhere to web standards and offer its best-performing web browser ever. Now, if you understand the company's history, you know that neither of these goals are particularly high bars. But IE 9, as it turns out, is quite impressive.

One of the weird side issues with web standards is that the various technologies—HTML 5, Cascading Style Sheets (CSS), JavaScript/ECMAScript, and so on—are in a constant state of flux. Various browser makers make claims about HTML 5 compliance, for example, despite the fact that the spec is in ongoing development, and is changing, and will be for a few more years at least. The W3C web standards body hasn't really shown up with too many industry-standard HTML 5 tests, but when they finally did so in late 2010, surprise, surprise, Microsoft won. And by a long shot.

Of course, the anti-Microsoft crowd doesn't rely on the W3C to prove which products are superior. They've created non-standardized tests, like the SunSpider JavaScript rendering suite—which doesn't measure real-world browser performance—and the ACID3 HTML 5 test, which doesn't measure HTML 5 standards compliance but other things, including some HTML 5 features that are already known to be changing. IE has never performed well in either test.

Until IE 9, that is. IE 9 will never achieve a perfect score in ACID3, because Microsoft refuses to support features in its browser that will be changing. But it does score a very respectable 95 out of 100, on those HTML 5 features that aren't guaranteed to change. IE 9 also scored the highest-ever rating on the SunSpider test, which must have been a shock to Google and Apple, given their incessant touting of their own products' superiority in JavaScript execution.

What I appreciate about all this is that Microsoft recognizes the futility of it all. These tests mean nothing in the real world, but they are always held against IE. So the company is doing well on the tests that matter—even though they don't really matter, if you get my meaning—and focusing too on real-world performance. Here, the rating is a bit more subjective. I find IE 9 to be on par with browsers like Chrome or Firefox from a performance perspective, using websites that are currently popular or typical.

Where IE 9 should really pull ahead, however, is in the coming generation of content- and feature-rich websites. That's because IE 9 offers something that no other browser will ever offer on Windows: Complete hardware acceleration. Yes, Chrome, Firefox, and even Safari will offer some forms of hardware acceleration, usually for certain content types only. But only IE 9 will be accelerated across the board. It's a huge advantage, one that makes websites run like native Windows apps, and one that points to the future, if you will, of hybrid web apps on Windows.

But I suspect we'll need to wait for Windows 8 before that vision is realized.

Windows Phone and the Carriers

One of the biggest promises of Windows Phone 7 was that Microsoft would bypass the carriers completely and deliver software updates both major and minor directly to users, just like Apple does. It was a wonderful promise. It's also completely untrue.

As it turns out, Microsoft's wireless carrier partners have the ability to throw a wrench into the gears of progress and prevent users from getting a software update from Microsoft. And this ability is explicitly provided to them by Microsoft as a weird concession of sorts, even though the software giant could simply choose to deliver updates through the Zune PC software and bypass the carrier networks altogether.

However, Microsoft will be testing each update against the carriers' own internal tests and will provide carriers with those results, proving that the updates won't harm users' phones. The carriers will likely let these updates simply sail through.

However, if a carrier feels that an update warrants more testing, it can prevent the update from appearing on its users' devices. The good news is that Microsoft's updates are cumulative, and carriers can only prevent an update from appearing until the next update is made available. So the blocking capability is temporary at best.

Still, I'm curious to see how aggressive carriers get after the first-generation Windows Phone buyers start getting close to the end of their two-year subscription term. I believe that these guys will put the brakes on and try to get customers to purchase a new phone instead of getting more free updates for their old one. Mark my words: We can't trust these companies.

Kinect and the Future of Computer UIs

Microsoft's Kinect motion sensor add-on for the Xbox 360 doesn't seem like the type of thing that would keep admins and IT pros up at night, but there is growing evidence that the software giant intends to add this interaction model to desktop versions of Windows. This makes some sense. We've been stuck in a fairly rote interaction model for decades now, and even touch and multi-touch interfaces from more recent Windows versions haven't done much to change that.

The reason Kinect is perhaps more important is that it signals a change to pervasive computing. This is a somewhat abstract concept, but I think of it as the difference between interacting with a thing—a tablet, laptop, or whatever—and interacting with your entire environment. When you type at a keyboard, you're doing a very specific action. But with some future version of Kinect, you could be silently and seamlessly triggering events on computers across the world as you move around in a room.

Of course, PCs being PCs, Kinect isn't really going to replace anything else—it's going to augment the other interaction models. So I suspect we'll be mousing around and using keyboards for years to come. But it will be interesting to see how motion sensors change everything, not just PCs but house and office designs, cars, and more. It's the future, and if you want to get started, I guess a quick game of "Kinect Adventures" isn't such a horrible way to get attuned to it.

InstantDoc ID 129032

PAUL THURROTT (thurrott@windowsitpro.com) is the senior technical analyst for *Windows IT Pro*. He writes a weekly editorial for *Windows IT Pro UPDATE* (www.windowsitpro.com/email) and a daily Windows news and information newsletter called *WinInfo Daily UPDATE* (www.wininformant.com).

Smarter technology for a Smarter Planet:

What database integration means to this blood sample.

It means doctors in Ethiopia will be able to instantly compare this blood sample to over 41,000 HIV treatment histories to help their patients receive the best treatment regimen possible. The EuResist Network is helping doctors predict patient response to various HIV treatments with over 78% accuracy—outperforming 9 out of 10 human experts in a recent study. The tool is built on an IBM analytics solution that integrates a variety of disparate databases onto a flexible IBM DB2® platform to process complex metadata more effectively than anything else on the market. A smarter organization is built on smarter software, systems and services.

Let's build a smarter planet. ibm.com/hospital

IBM®

A data visualization of 41,000 HIV case histories.

Minasi

"Save time and space by adding /delete and /export to your toolkit."

Managing Multiple-Image Files with ImageX

Two essential options increase your imaging capabilities

In my past few columns, I've been showing you how to use ImageX. Last month, I showed you the ImageX /append command, the key to creating multiple-image files. This month, I want to take that discussion a bit further and show you how to use ImageX to *manage* multiple-image .wims.

How do you determine whether a .wim file contains more than one image? And how do you find out what's in those images? You use the /info option, as in

```
imageX /info images.wim
```

ImageX returns some information as text, including the number of images in the .wim file. The rest is in simple XML, so you can easily find overall size information about each image (<DIRCOUNT>, <FILECOUNT>, and <TOTALBYTES>) and the image's descriptive name (<NAME>). The descriptive name is the one ImageX requires you to specify when you use the /capture or /append switches.

For a closer look at an image, you'd have to switch tools and use the Dism command with the /get-wiminfo option:

```
dism /get-wiminfo /wimfile:<filename> /index:<image-number>
```

Dism's output goes beyond byte and folder numbers and reports on the OS inside the .wim file, including SKU and service pack level.

Beyond viewing image information within a .wim file, ImageX lets you remove an image from a .wim altogether with the /delete switch. To remove an image from a .wim, just type

```
imageX /delete <sourcewimfilename> <imagenumber|imagename>
  [/check]
```

Suppose you have a file called images.wim that contains two images, and the second image's name is *EnterpriseBuild*. You could remove that second image by typing either

```
imageX /delete images.wim 2
```

or

```
imageX /delete images.wim enterprisebuild
```

Note that image names aren't case-sensitive, wild cards don't work, and you can't use /delete to remove the only remaining image from a .wim file. Also, note that if you created the original images with the /check option (as I discussed a few columns ago), you should also use /check in your delete operation to preserve the hashes.

The .wim file's ability to hold multiple images in a single file while being stingy with your disk space is nice. But sometimes you'll want to pull out just a single image and put it in another .wim file. In that case, you'd use the /export switch:

```
imageX /export <sourcewimfilename> <imagenumber |
  imagename> | * <destinationwimfilename> ["<new image
  name>"] [/check]
```

For example, to export the third image—an image named *win7kiosk*—from a file named images.wim into a file named justthree.wim, you could type

```
imageX /export images.wim 3 justthree.wim
```

That command works whether justthree.wim exists or not. As with /delete, you could have replaced the *3* with *win7kiosk* to the same effect. The /export option is unlike /delete, however, in that it accepts the wildcard asterisk character, and in that case would export every image in images.wim into the (now sadly misnamed) justthree.wim file.

When does an asterisk make sense? Consider a case in which you have several .wim files, each of which contains a bunch of images, and you want to consolidate all your .wims into a single .wim. To accomplish that, you'd need only execute—for each of your .wim files—the command

```
imageX /export <filename> * masterlibrary.wim
```

The /export option lets you copy images from one .wim file into a new .wim file. I wish Microsoft had chosen the option name /copy rather than /export because I have trouble remembering whether /export deletes an image from the source .wim file after exporting it. It doesn't. If you want the exported copy of the image to have a different name than its original one, just add one in quotes after the name of destination .wim in the /export command, as in

```
imageX /export images.wim 3 justthree.wim
  "locked-down image"
```

InstantDoc ID 128928

MARK MINASI (www.minasi.com/gethelp) is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 25 books.

Otey

"Host-level backups can't be considered a substitute for guest backups. Applications need to be backed up at the guest level for end-user data protection."

Virtualization Mistakes

Avoid these security, resource, and deployment errors in your virtual environments

Virtualization has become a core infrastructure technology for most businesses. It's so prevalent that many organizations plan for virtualization of all new server deployments. One factor that has allowed virtualization to become so popular is its ease of implementation. However, that easy setup can come back to bite you if you don't plan your virtual deployments properly. Let's look at the top ten virtualization mistakes you should avoid.

10 **Virtualizing on older hardware—**Both Microsoft Hyper-V and VMware's ESX Server can run on older hardware platforms. However, newer processors have features such as Second Level Address Translation (SLAT) and Nested Page Tables (NPT) that can drastically improve virtualization performance by letting the hardware take care of the translation between the guest virtual machine (VM) memory addresses and the physical RAM addresses.

9 **Running antivirus on virtual hard disks—**Implementing antivirus protection is always a good idea. However, antivirus scans on a VM's virtual hard disk can degrade the performance of the VM. Be sure to exclude virtual hard disks from the host's antivirus scanning.

8 **Ignoring guest VM backup—**You can back up the VM at the host level without interrupting end-user services, enabling easy disaster recovery because you can quickly restore that host image on another virtualization host. Even so, host-level backups can't be considered a substitute for guest backups. Applications such as Microsoft SQL Server and SharePoint need to be backed up at the guest level for end-user data protection.

7 **Inadequate virtualization host security—**It's easy to focus on guest security, but securing the host is even more important because the host has access to all guest resources. Hosts must have physical security, plus all the resources on the host should be secured according to the principles of least privilege.

6 **Always using the VM default settings—**Another common mistake is to blindly accept the default settings used by the virtualization host and the VM management console. I commonly change the default VM location from DAS to SAN. In addition, you need to take care to size each VM's CPU, RAM, disk, and network to match the workload required by that VM.

5 **Inadequate host processor resources—**Virtualization lets you achieve much higher hardware utilization rates than you can reach with a physical server. However, nothing stops you from overcommitting your host CPU with too many virtual workloads. Ideally you want to have one host CPU core per VM. Windows Server Resource Monitor can give you a quick heads up about your CPU and core utilization levels.

4 **Inadequate host RAM—**RAM is the primary limiting factor to the number of concurrently active VMs because each VM must allocate its RAM from physical memory. Make sure you have an adequate amount of host RAM for the number of VMs you plan to run as well as leaving enough RAM for the virtualization host.

3 **Inadequate host network adapter cards—**Another common mistake—especially in server consolidation projects—is failing to install an adequate number of network adapter cards in the virtualization host. In a server consolidation environment, all of the network bandwidth from the VMs is funneled through the host's network adapters. Although you might not need a one-to-one relationship, it's easy to overcommit a few network adapters with the traffic from many VMs.

2 **Too many VMs per Cluster Shared Volume—**Cluster Shared Volumes (CSVs) are a new feature in Windows Server 2008 that lets multiple VMs share the same LUN. By default, all VMs go to the same CSV. This design can be OK for light workloads, but heavier workloads, such as SQL Server, require more CSVs. In addition, remember that disk performance depends on the number of spindles, so using storage with a large number of spindles provides better performance.

1 **Using only one CSV per VM—**It seems common to think that a VM is limited to using one CSV. Not only can you create more than one CSV per virtual server, but you can also split your VM's VHD files between CSVs. You could split your system files and page file to a VHD on one CSV volume and put your data files and user data in a VHD file located on another CSV volume.

InstantDoc ID 129009

MICHAEL OTEY (motey@windowsitpro.com) is senior technical director for *Windows IT Pro* and *SQL Server Magazine* and author of *Microsoft SQL Server 2008 New Features* (Osborne/McGraw-Hill).

Yet Another 10 Free Tools for System Administrators

Audit Active Directory and file servers, detect inactive users, block USB devices, and more – for free

The following freeware tools by Windows IT Pro Community Choice Awards finalist NetWrix Corporation can save you a lot of time and make your network more efficient – at absolutely no cost. Some of these tools have advanced commercial versions with additional features, but none of them will expire and stop working when you urgently need them.

1. Active Directory Change Reporter (Windows IT Pro Sep'09: InstantDoc ID 102446, Windows IT Pro Jan'09: InstantDoc ID 100593, TechTarget: www.tinyurl.com/2ulfzy6) — This is a simple auditing tool to keep tabs on what's going on inside Active Directory. The tool tracks changes to users, groups, OUs, and other types of AD objects, and sends summary reports with full lists of what was changed and how it was changed. In addition, it has a nice "rollback" feature that helps rollback unwanted changes (including deletions) very quickly. Download link: www.tinyurl.com/3xgmcsc

2. USB Blocker (Windows IT Pro Nov'09: InstantDoc ID 102860) — Users bring tons of consumer devices: flash drives, MP3 players, cell phones, etc., into the office and this aptly-named tool can block them with a couple of mouse clicks to prevent the spread of a virus and to restrict the take-out of confidential information. The product is integrated with Active Directory and is very easy to use. Download link: www.tinyurl.com/2vhxuyc

3. Password Expiration Notifier (Redmond Magazine Feb'09, 4sysops: www.tinyurl.com/35w4xhw) — This tool will automatically remind users to change passwords before they expire to keep you safe from password reset calls. It works nicely for users who don't log on interactively and, thus, never receive standard password change reminders at log on time (e.g., VPN and OWA users). Download: www.tinyurl.com/373apnm

4. Inactive Users Tracker (MS TechNet Magazine May'08: www.tinyurl.com/3a6sxom) — This feature tracks down inactive user accounts (e.g., terminated employees) so you can easily disable them, or even remove them entirely, to eliminate potential security holes. The tool sends reports on a regular schedule, showing what accounts have been inactive for a configurable period of time (e.g., 2 months). Download link: www.tinyurl.com/2uk3dj6

5. File Server Change Reporter (4sysops.com: www.tinyurl.com/38zmx55) — This tool enhances the line of auditing tools; this one for file servers. File Server Change Reporter detects changes in files, folders, permissions, tracks deleted, and newly-created files, and sends daily summary reports. This is a very useful tool to detect mistakenly-deleted files and recover from backup or to see if someone changes some important files. Download link: www.tinyurl.com/2utcprx

6. Active Directory Object Restore Wizard (4sysops.com: http://tinyurl.com/3ys5sq2) — This tool can save the day if someone accidentally (or intentionally) deleted a bunch of Active Directory objects. It provides granular object-level and even attribute-level restore capabilities to quickly rollback unwanted changes (e.g., mistakenly deleted users, modified group memberships, etc). Download link: www.tinyurl.com/2vqzg2u

7. VMware Change Reporter (TechTarget/SearchVirtualDesktop: www.tinyurl.com/39emdlb) — If you don't know what is being changed by your colleagues in the VMware infrastructure, it's very easy to get lost and miss changes that can affect the things for which you are responsible. This tool tracks and reports configuration changes in VMware Virtual Center settings and permissions. Download link: www.tinyurl.com/39wyg6q

8. Windows Service Monitor (WindowsReference.com: www.tinyurl.com/2vslat4) —This very simple monitoring tool alerts you when some Windows service accidentally stops on one of your servers. The tool also detects services that fail to start at boot time, which sometimes happens, for example, with Exchange Server.
Download link: www.tinyurl.com/39h9j4w

9. Bulk Password Reset (reviewed by SoftPedia: www.tinyurl.com/39kaex5) – While most companies have strong password policies for their employees, one critical issue is still neglected: local Administrator passwords on all servers are usually managed in a "set and forget" fashion, sometimes using some "well-known" passwords, opening a major surface for security attacks. The Bulk Password Reset tool quickly resets local account passwords on all servers at once, making them more secure. Download link: www.tinyurl.com/35od5cv

10. Disk Space Monitor (MS TechNet Magazine Sep'09: www.tinyurl.com/2v44z2x) — Even with today's terabyte-large hard drives, server disk space tends to run out quickly and unexpectedly. This simple monitoring tool will send you daily reports regarding all servers that are running low on disk space, below the configurable threshold.
Download link: www.tinyurl.com/32ew3ep

“THE CONVERSATION BEGINS HERE”

UNIFIED COMMUNICATIONS CONNECTIONS

Microsoft® Exchange CONNECTIONS

WINDOWS CONNECTIONS

SharePoint CONNECTIONS

Microsoft® Visual Studio® CONNECTIONS

Microsoft® ASP.net CONNECTIONS

Microsoft® Silverlight® CONNECTIONS

SQL Server CONNECTIONS

BONUS: Mobile Apps Track

QUESTIONS ANSWERED • STRATEGY DEFINED • RELATIONSHIPS BUILT

Make **CONNECTIONS** the **CONFERENCE** you bring your whole team to this year!

One Place, One Time...

MARCH 27-30, 2011

ORLANDO, FL

GRANDE LAKES JW MARRIOTT RESORT HOTEL

Only Microsoft and Industry Experts speak at WinConnections!

WinConnections ... Providing the vision + intelligence to keep you and your company competitive in today's market!

KEYNOTES AND INDUSTRY EXPERTS

QUENTIN CLARK
MICROSOFT

STEVE FOX
MICROSOFT

SCOTT GUTHRIE
MICROSOFT

TIM HUCKABY
INTERKNOWLOGY

TONY REDMOND
TONY REDMOND AND ASSOCIATES

PAUL THURROTT
WINDOWS IT PRO

CHECK WEB SITE FOR DESCRIPTIONS OF SESSIONS AND WORKSHOPS

www.WinConnections.com • 800.505.1201 • 203.400.6121 • Register Today!

Microsoft · SharePoint Pro CONNECTIONS · SQL Server · Windows IT Pro · TECH Conferences Inc. PENTON MEDIA

Deuby

"Despite worries about virtualizing and moving applications wholesale to the cloud, identity stores aren't going anywhere in the forseeable future."

Should Identity Professionals Fear the Cloud?

Cloud computing has shaken up your IT department, demanding that you build an entirely new skill set

Welcome to my new column! I've been writing for *Windows IT Pro* nearly since its inception, and now I'm very pleased to say I've joined Penton Media as Technical Director for *Windows IT Pro* and *SQL Server Magazine*, as well as the Connections conferences. In this new role, I'll be involved in cloud computing and Windows Server topics, while continuing my focus on Active Directory (AD). I also get to write this new column, which will explore the range of technologies, concerns, and opportunities that identity professionals face today. This month, I want to talk about the evolution of enterprise identity management and cloud computing's effect on it.

Cloud computing—hot on the heels of virtualization—has shaken up the IT department. Unlike the move to virtualization (which at least initially remains within IT's data centers), cloud computing solutions can allow IT's traditional customers to perform an end run around established IT processes. Departments can buy Software as a Service (SaaS) applications such as Salesforce.com or Google Apps to provide capabilities that don't require any capital expenditure, scale well, and have setup times measured in hours or days rather than weeks or months.

> The biggest milestone in the history of Windows Server identity management was the introduction of Windows 2000 and AD.

Cloud computing has also shaken up another aspect of IT: the IT professional's career path. What kind of a future does the IT pro have if computing is evolving from on-premises data centers to massive but invisible data centers in the sky? There are a lot of aspects to this question, but I want to focus on a career path I'm most familiar with: the identity professional.

I've been working with Windows identity management in one form or another for many years, and I've watched it evolve over that time. For its entire existence, the scope of responsibility has grown—never shrunk—and I don't think that will change in the near future. In the beginning, there was the standalone computer, which had only a local account database. If you wanted to access resources on another computer, you had to create an account on it. Workgroup computing, introduced with Windows for Workgroups (WFW) and carried forward to this day, has separate account databases on each system, but you can authorize a user on another system to access your resources.

LAN Manager in the early 1990s marked the transition of account management from PCs to servers. It was the first Microsoft networking product to have a single user-accounts database server that determined identity and access across multiple servers, and to authenticate clients to those server's resources. It was a pretty cantankerous piece of software, put together from OS/2, MS-DOS, and 3Com's 3+Share network server software. I remember going into the morning operations meeting at Texas Instruments, where we'd beat up the Microsoft team over the previous day's outages.

In 1993, Microsoft developed what we OS geeks considered the real thing—Windows NT 3.1, a 32-bit pre-emptive, multiprocessing, multithreaded OS that was far more capable than its predecessors. With this leap in capability came the NT domain, a much more capable account-management system. Instead of a single account server, an NT domain had a primary *domain controller* (DC) and multiple backup DCs to ensure that the domain would keep functioning if the primary DC was offline.

NT domains were orders of magnitude more scalable than LAN Manager; they could hold several thousand user accounts and computers. Even better, you could configure an NT domain to hold only accounts (account domains) or only computers (resource domains) and string these together so that the computers in the resource domains used the accounts in the account domains to control access to their data. And you could string many of these domains together with trust relationships to handle many thousands of users and computers. This scenario spurred IT's first attempts to wrangle user account management out of the hands of the departments and into a centrally managed organization. Manual account management is a thankless job, and departments in general were happy to hand the work to someone else. This got very complicated, though, as the number of account domains and their trusts multiplied dramatically. For example, Intel had 156 trusts between its account domains alone, and many hundreds

more between them and the resource domains.

As the Windows network increased its scope to a significant segment of the enterprise, it began to show up on high-level management dashboards; for many companies, it became the primary network authentication mechanism. Even so, if you'd asked an NT administrator whether he or she was an identity manager, you'd have received a blank stare. But at this scale, much of an admin's time was taken up with managing a user's identity and access management across these many domains.

The biggest milestone in the history of the Windows Server OS, and its identity management, was the introduction of Windows 2000 and AD. Designed to scale to a level that still hasn't seriously challenged its limits 10 years after it was introduced, a single AD forest of domains was first thought to be the only solution you'd ever need—*one forest to rule them all!* At Intel, through some very good negotiations, we managed to convince our various business units to do just this, leaving behind their NT 4.0 fiefdoms and moving to organization units (OUs) in a single corporate forest. Once again, this increased the scope of Windows identity management. Soon, however, many companies realized that Windows security depended not on domain boundaries but on forest boundaries, and multiple forests started popping up. Business needs often dictated that these forests needed to share some information, so inter-forest trusts were set up. As a result, some companies' AD environment began looking like a scaled-up version of their NT 4.0 environment, with multiple AD forests taking the place of multiple NT domains.

As the scope of Windows identity management continued to grow, technical decision-makers and AD administrators began thinking about it in a larger context. Instead of thinking about identity from only within the Windows network ("How do we manage user accounts and passwords within our AD forests?"), they began viewing it as part of the enterprise ("How do we manage the employee's identity information holistically across the entire company?"). Though by now our AD credentials controlled most of our access to the company's computing infrastructure, there were still significant gaps—most notably, human resource databases and physical security systems. These databases contained sensitive and important aspects of an employee's digital identity, but they were loosely coupled by manual batch jobs if they were linked at all. This caused serious security concerns—for example, in the employee-termination scenario: If all of a user's accounts aren't disabled immediately throughout all systems he or she has an account on, the terminated employee might be able to cause damage to a system that still contains the user's active account. Identity professionals began using metadirectory services such as Microsoft Identity Lifecycle Manager (ILM) and its successor Forefront Identity Manager (FIM) to take these disparate identity sources and link them in a metadirectory (literally, *directory of directories*) that provides a holistic view of everything associated with a user's identity, and the ability to shape and control data flows between these identity sources. With a metadirectory service overseeing all the company's identity sources, a termination action triggered in the HR database cascades into other systems, immediately locking out Windows accounts and any other systems the user had access to (e.g., disabling badge access to buildings). Implementing metadirectory services isn't a trivial project, however: The software is costly, the learning curve is pretty steep, and the implementation is complex. As a result, only a relatively small percentage of AD professionals have added this technology to their skill set so far.

> Cloud identity is the next step in the natural progression of growth for the identity professional.

Now, we have cloud computing and its identity-management concerns to contend with. Cloud identity—the management of identity credentials between identity providers (e.g., your AD forest) and cloud service providers (e.g., Amazon's Elastic Computer Cloud)—is the next step in this natural progression of growth for the identity professional. By default, the scope of your corporate identity system doesn't extend to the cloud. The challenge is how to securely use your existing corporate identity for these services. We've all experienced the dizzying number of user accounts and passwords necessary for consumer websites; why should you, similarly, have to create and manage a separate set of accounts for every cloud service provider you use at work? Readying your company to work securely with cloud services is relatively cheap, though the technology and its jargon can be confusing.

Federation is the leading technology for stitching your identity system to the cloud service of your choice. It lets you securely provide just the relevant identity data the cloud application needs without exposing the entire identity system externally. Microsoft's Active Directory Federation Services (AD FS) and Ping Identity's PingFederate are examples of established federation products available today. Understanding federation and the claims-based authentication model it supports is the next step for the identity professional who wants to keep his or her skills current because we're at the beginning of a growth curve for this technology. (I'll be covering cloud identity technologies such as federation in future columns.)

Cloud computing isn't going away, but that's a good thing for identity pros. Despite worries about virtualizing and moving applications wholesale to the cloud, identity stores aren't going anywhere in the foreseeable future. Basic security principles dictate that AD stays with you, on the premises. Cloud computing gives you a new challenge and skill set to acquire. The cloud isn't a threat to identity professional's careers but rather an opportunity.

If you have a topic you'd like to see discussed here, please drop me a line! This magazine is all about reflecting what's on the mind of the IT professional. See you next month!

InstantDoc ID 129026

SEAN DEUBY (sean@windowsitpro.com) is technical director for *Windows IT Pro* and *SQL Server Magazine,* and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.

- Certification Authorities
- Hyper-V
- Superfetch
- Internet Explorer
- VMware

ANSWERS TO YOUR QUESTIONS

Q: Is there an easy way to automatically re-enroll certificate holders that received a certificate from an old CA with a new certificate in a new PKI hierarchy?

A: To force all holders of a particular certificate to automatically enroll for a replacement certificate issued by a CA in your new PKI hierarchy, use the Reenroll all Certificate Holders feature of the Certificate Templates MMC snap-in. All you need to do is right-click the certificate templates you want to reenroll and select Reenroll All Certificate Holders from the context menu.

Behind the scenes, this action will increment the version number of the certificate templates. This change is then detected by the auto-enrollment service on your Windows workstations and servers that will trigger your users and computers to enroll for certificates of the updated templates.

The auto-enrollment service updates user and computer certificates at the next auto-enrollment pulse. For computers, the auto-enrollment pulse occurs at computer startup and every eight hours. For users, the auto-enrollment pulse occurs at user logon and every eight hours. You can also manually trigger an auto-enrollment pulse by running the following command from the command line:

```
certutil -pulse
```

Certutil.exe is included in the Windows Server 2003 Administrative Tools. In later Windows versions, this tool is installed by default.

For this automatic certificate re-enrollment to work, you must also make sure that a specific GPO setting is enabled. In the GPO auto-enrollment properties, you must have the Update certificates that use certificate templates option selected.

The auto-enrollment properties are located in the Computer Configuration\Windows Settings\Security Settings\Public Key Policies and User Configuration\Windows Settings\Security Settings\Public Key Policies GPO containers. You must enable this option in both the user and computer configuration auto-enrollment properties to allow administrators to force both computers and users to reenroll for an updated certificate template.

—Jan De Clercq

InstantDoc ID 128986

Q: When should I use NAT networking with a virtual machine (VM) in VMware Workstation?

A: You probably configure your VMware Workstation VMs to use Bridged networking. With bridged networking, the virtual network switch created by VMware Workstation bridges whatever network the host uses to the VM. By doing this, whatever IP address range used by your host is also used by your VMs.

This configuration works great if your host doesn't move around a lot, but it can cause problems when your host's subnet changes. This becomes a very big problem when you've assigned static IP addresses to VMs.

You can get around this problem by connecting VMs to Workstation's NAT interface (VMnet8) instead of its bridged network (VMnet0). VMs connected to this network are given a special subnet—the actual subnet is randomly generated the first time you launch the Virtual Network Editor. That subnet only changes when you specifically configure it to. VMs in that network needn't suffer a subnet change when the host does.

Be aware that NAT comes with its own limitations. The host will always be able to connect to VMs on the NAT network, but other computers won't. If you need other computers to connect to VMs on the NAT network, create port forwarding rules under NAT Settings. These rules, which are set by the individual TCP or UDP port, will let external computers talk with VMs in the NAT network.

—Greg Shields

InstantDoc ID 128789

Q: I have a 64-core machine that has hyperthreading, for a total of 128 logical processors. When I enable the Hyper-V role, I only see 64 logical processors. Why?

A: Hyper-V only supports 64 logical processors, so once the Hyper-V role is enabled, only the first 64 logical processors will be exposed. If you're using hyperthreading on your CPUs, this would mean the first 32 physical cores and the 32 hyperthreaded "cores." Because a physical core is preferable to a hyperthreaded one, it's recommended that if you have 64 physical cores, disable hyperthreading. That way, all 64 physical cores will be available to Hyper-V and you'll get the best performance possible.

—John Savill

InstantDoc ID 128962

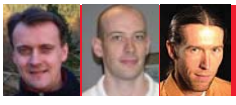

Jan De Clercq | jan.declercq@hp.com
John Savill | jsavill@windowsitpro.com
Greg Shields | virtualgreg@concentratedtech.com

Q: What's the length of the grace period where I can run a Windows Server 2008 R2 Remote Desktop Services (RDS) host without it requiring valid RDS CALs from a licensing server?

A: When the RDS Host role is enabled on a server, many user connections to the server are possible, providing an RDS CAL can be assigned to the user or device (depending on configuration) from a licensing server. A newly deployed RDS host can accept connections for a grace period of 120 days before a valid CAL must be available from a licensing server. This 120-day grace period is the same for Windows Server 2008, 2003 R2, and 2003. Windows 2000 had a grace period of 90 days.

—John Savill

InstantDoc ID 128888

Q: How can I make it so that non-primary IP addresses on a network adapter don't register in DNS with Windows Server 2008 R2?

A: In Windows Server 2008 R2, the behavior of the non-primary addresses on an adapter changed so that all IP addresses register with DNS instead of just the primary IP address. Microsoft released a hotfix for Windows 7 and Windows 2008 R2 that allows you to configure additional IP addresses on an adapter without them registering with DNS. It uses the skipassource=true flag. You need to:

1. Apply the hotfix.
2. Remove the additional IP addresses from the network adapter.
3. Add the IP addresses back using netsh. For example,

```
netsh int ipv4 add address
   &lt;interface&gt; &lt;ip
```

—John Savill

InstantDoc ID 128981

Q: Does it make sense to defrag Windows computers that are VMware virtual machines (VMs)?

A: Reader Blue Michael Planté of La Habra, California asked the question, "I cannot get a concrete answer. Does it make sense and should we defrag our Windows machines that are in a VMware VM?"

The answer has two parts, because fragmentation in VMware virtual environments can actually happen in two different places. First is within the VMFS file system where your VM's VMDK (and other) files are located. The VMFS file system is substantially different from the NTFS file system in Windows. Among other traits, its block size is significantly larger, with VMFS' block size measured in megabytes as compared to NTFS' kilobytes. Files inside the VMDK file system also tend to be larger than files in NTFS. While fragmentation can and will happen inside the VMDK file system, these two facts tend to reduce its effect on performance.

Thick disks, those that are pre-allocated as they're created, also tend to have less fragmentation within the VMDK file system than thin-provisioned disks. Because the entire size of a thick disk VMDK file is allocated at once, these disks tend not to grow fragmented. Thin disks, because they can grow over time, will experience slightly more fragmentation.

In either case, most environments needn't worry about fragmentation's impact on performance within the VMFS file system. VMware's knowledgebase article 1006810 discusses this conclusion in more detail.

Conversely, the NTFS file system inside your Windows VMs does experience performance-affecting fragmentation. This fragmentation occurs whether the Windows computer is physical or virtual. As a result, you probably should defragment the disk files of Windows computers.

One element of defragmentation utilities running inside VMs warrants attention. The resource utilization of a defragmentation utility can be significant. Running the defragmenter can spike resource use, taking away resources from other VMs on the host—the defragmentation process itself could have an effect on VM performance.

Be careful not to enable a defragmentation pass on every VM at once, or you might severely hurt performance. Also, consider exploring one of the new generation of virtualization-aware defragmenters that's on the market today.

—Greg Shields

InstantDoc ID 128989

Q: How can I remove Internet Explorer (IE) from my Windows Server 2008 R2 (full install) server?

A: If you need to remove IE from a Server 2008 R2 full install server (Server Core doesn't have IE installed), you can use the DISM command. I stress the IE can be very useful and if you just want to stop users from accessing the Internet or stop certain people from running IE, Group Policy might be a better option. You list all your features using

```
dism /online /get-features
```

One of the features is Internet-Explorer-Optional-amd64, which you can remove with:

```
dism /online /disable-feature /
   featurename:Internet-Explorer-
   Optional-amd64
```

After you restart, it won't be available. If you need to add it back, replace /disable-feature with /enable-feature in the command above.

—John Savill

InstantDoc ID 128963

Q: How can I hide certain drive letters from Windows Explorer and My Computer?

A: Microsoft provides a Group Policy Object (GPO) setting to hide certain drive letters from My Computer and the Windows Explorer. Hide these specified drives in My Computer is located in the User Configuration, Administrative Templates, Windows Components, Windows Explorer GPO container. The setting also affects the drive letters that appear in the standard Windows Open dialog box.

To use this setting, navigate to the above GPO container in the Group Policy Editor MMC snap-in (gpedit.msc), open the setting, and select the Enabled radio button. In the dropdown box, you can then select what specific drive or combination of drives you want to hide.

This setting only removes the drive icons from Windows Explorer, My Computer, and Windows Open dialog boxes. Users can still gain access to the contents

of the hidden drives using other methods, such as typing the path to a directory on the drive in the Map Network Drive dialog box, using the Run dialog box, or using a command window. This setting also doesn't prevent users from using programs to access these drives or their contents, and doesn't prevent users from using the Disk Management snap-in to view and change drive characteristics.

More information on this setting can be found in this Microsoft Knowledge Base article. This article also explains how you can modify the GPO administrative template files (*.adm) to display your custom list of drives or collection of drives in the dropdown box of the above GPO setting.

A related setting located in the same GPO container is Prevent access to drives from My Computer. This setting can prevent users from using My Computer to access the content of selected drives. If you enable this setting, users can still browse the directory structure of the selected drives in My Computer or Windows Explorer, but they can't open folders and access the contents of the drives. Also, they can't use the Run dialog box or the Map Network Drive dialog box to view the directories on the drives.

—Jan De Clercq

InstantDoc ID 128783

Q: Should I disable Superfetch if I have a solid state disk (SSD)?

A: Superfetch is a great feature that preloads programs into memory. When the user launches the programs, they start faster than normal. However, if you have an SSD drive, random reads are very fast. So does Superfetch help or just cause more I/O hit to the SSD, decreasing its lifetime?

You can manually disable the Superfetch service by disabling it, but Windows 7 actually takes care of this for you. Supefetch checks the WinSAT disk score of the system (which considers random reads and writes) and if it's over 6.5, Superfetch automatically turns itself off.

Defrag performs a similar check and doesn't run if the device doesn't have a seek penalty—it checks IOCTL_STORAGE_QUERY_PROPERTY:StorageDeviceSeekPenaltyProperty.

If the system can't figure that out, it won't run defrag if your disk's random read rate is greater than 8MB/second.

—John Savill

InstantDoc ID 128954

Q: How can I Use VMware Disk Mount to interact with a VMDK File?

A: I use VMware Workstation for all kinds of software evaluation. I also use it while I'm on the road, teaching at conferences or consulting for my company, Concentrated Technology. Workstation is great when you plan to use it for a while, but the overhead of getting it started and powering on virtual machines (VMs) can be painful for simple file work.

For those situations where you just want to add, remove, or modify files in a .VMDK file, download a copy of the VMware Disk Development Kit. Among other tools, the kit includes a handy command line utility called vmware-mount, also known as VMware Disk Mount. You'll find the utility in C:\Program Files\VMware\VMware Virtual Disk Development Kit\bin.

Once it's mounted, you can work with that disk in Explorer, just like any other disk. To mount a local .VMDK to the M: drive, use the command

```
vmware-mount M: {pathToVMDKFile}
```

You can even use this tool to mount remote .VMDKs, either on other Windows hosts or ESX/ESXi hosts. Here's some quick syntax to connect to a disk on a remote ESX/ESXi host:

```
vmware-mount K: "[storage1] WinXP/
   WinXP.vmdk" /i:ha-datacenter/vm/
   WinXP /h:esx3 /u:root /s:secret
```

You can get all the command line hints from the tool's documentation.

Note that vmware-mount in the kit's 1.2 version doesn't work on x64 computers. That's a big omission that will hopefully be resolved in its next version. Also, this command won't work with disks when their VMs are running.

—Greg Shields

InstantDoc ID 128791

Q: I'm planning to set up an additional Windows enterprise Certification Authority (CA) in my Windows Server 2008 Active Directory (AD) forest. How can I make sure that the new CA doesn't start issuing certificates until I've properly and completely configured the CA?

A: You can block your new CA from issuing new certificates by ensuring that it doesn't have any preconfigured certificate templates when it boots up for the first time. You can accomplish this by adding the following line to the capolicy.inf configuration file:

```
LoadDefaultTemplates=False
```

This line must be added to the [certsrv_server] section of the capolicy.inf file. The LoadDefaultTemplates line controls whether the CA is configured with any of the default certificate templates. In the Windows CA configuration, certificate templates determine what types of certificates a CA can issue.

The LoadDefaultTemplates entry only applies during the installation of an enterprise CA—it doesn't affect a standalone CA. In the Active Directory Certification Service (ADCS), an enterprise CA is an AD-integrated CA. On Windows Server 2003 and Windows Server 2003 R2, the LoadDefaultTemplates setting only applies to root enterprise CAs and is ignored on a subordinate enterprise CA. On Server 2008 and Server 2008 R2, the LoadDefaultTemplates setting applies to both root and subordinate enterprise CAs.

In case you're not familiar with the capolicy.inf file: It's a configuration file that contains various settings that are used when installing a Windows CA or when renewing the CA certificate. The capolicy.inf file isn't required to install ADCS with the default settings, but in many cases the default settings are insufficient. Once you've created a capolicy.inf file, you must copy it to the %systemroot% folder of your server before you install the CA or renew the CA certificate. More information on the syntax of the capolicy.inf file can be found online at Microsoft Technet or MSDN.

—Jan De Clercq

InstantDoc ID 128985

Q: My machine was in sleep mode and I stepped away. When I came back to it, the machine was awake. How can I tell what caused the machine to wake?

A: I had a similar problem, where I put my machine in sleep mode and came back to see it wide awake, and I had no clue why. The easiest way to check why a machine woke from sleep is to run the command

```
powercfg -lastwake
```

This will show the reason why the machine exited sleep state.

—John Savill

InstantDoc ID 128956

Q: I'm running Windows Update in my XP Mode virtual machine (VM). It says some updates have been hidden. Should I unhide them?

A: No. If you check for updates that are hidden by default in your XP Mode VM, you'll see it's actually Internet Explorer (IE) 8, which would replace Windows XP's default IE 6.

The update is hidden because many people use XP Mode both to run applications that won't work in Windows 7 and to use Internet Exporer 6 with plug-ins and code that won't work with IE 8 or later. Microsoft hides the IE 8 update by default to stop people from accidently replacing IE 6.

—John Savill

InstantDoc ID 128955

Q: Is there an easy way to block certain groups of users from browsing the web with Internet Explorer?

A: You could restrict the use of iexplore .exe, but that can cause problems with the system, including local browsing. The easiest way to stop people from browsing the Internet is to set user machines to use an Internet Explorer proxy that's invalid and prevent them changing the proxy configuration. You also need to ensure they can't install an alternate browser, such as Firefox, which would bypass your Internet Explorer proxy configuration.

1. Open a group policy object that applies to the users and move to User Configuration, Policies, Windows Settings, Internet Explorer Maintenance, Connection, Proxy Settings.
2. Enable the proxy settings option and set a dummy IP address for HTTP and Secure, such as 0.0.0.0, and click OK.
3. Go to User Configuration, Policies, Administrative Templates, Windows Components, Internet Explorer, Internet Control Panel.
4. Double-click Disable the Connections page and set it to Enabled. Now users can't change their proxy configurations.

—John Savill

InstantDoc ID 128957

Q: We want to add more reliability to our Windows Public Key Infrastructure (PKI). One way to provide this is by installing extra Certification Authority (CA) servers. Can a Windows CA also be clustered?

A: Yes, starting with Windows Server 2008, a Windows CA can be clustered. CA clusters are supported in the Windows Server 2008 Enterprise and Datacenter editions. You can only use clustering for the AD Certificate Services (CS) service and not for other AD CS role services, such as the Online Responder service (this is the service providing OCSP support) or the Network Device Enrollment Service (This is the service providing SCEP support).

Windows CAs can only be configured to use a two-node active/passive cluster. This means that the CA service is only active on one cluster node at a time. If the active node becomes unavailable, the second node becomes active and the CA service will resume on the second node. Windows Server 2008 does not support active/active CA clusters.

CA clusters require shared storage to make the CA database and log files available to both nodes. When you use Hardware Security Module (HSM) to protect your CA keys, you'll also need a shared HSM. This can only be provided if you use a network-attached HSM.

—Jan De Clercq

InstantDoc ID 128987

Q: Can I block machines from accessing a wireless network using Group Policy?

A: It's possible for Windows Vista and later clients.

Open an existing Group Policy Object or create a new one. Navigate to Computer Conifguration, Policies, Windows Settings, Security Settings, Wireless Network (IEEE 802.11) Policies. Right-click in the contents pane and select Create a New Wireless Network Policy for Windows Vista and Later Releases. Provide a name and description, then select the Network Permissions tab. Click the Add button to add a new wireless network policy.

Enter the SSID of the wireless network you want to block. Select the type and set the Permission to Deny and click OK to all dialog boxes. Ensure the GPO is linked to the OU/domain/site whose computers should not use the specified SSID.

—John Savill

InstantDoc ID 129080

Q: System Center Mobile Device Management is discontinued and its features will be rolled into System Center Configuration Manager (SCCM) 2007 R3. What licenses do I need to manage mobile devices?

A: SCCM has two types of management license for non-server OS environments—OSE management licenses (one OSE for each OS, which can be used by multiple users) and user management licenses (one user management license for each user, allowing all devices assigned to the user to be managed by SCCM).

If your users currently have SCCM user management licenses, their mobile devices can be managed by SCCM without additional licenses. If OSE management licenses are currently assigned to the desktops, additional OSE management licenses will be required for the mobile devices, unless you can switch to user management licenses. (This will also be the case if mobile devices need to be managed for users that don't have a SCCM-managed desktops.)

—John Savill

InstantDoc ID 129074

Microsoft Brings SBS into the Cloud

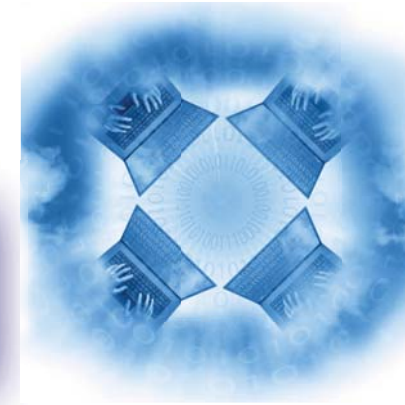

A new lower-end, cloud-based Essentials offering targets small businesses

by Paul Thurrott

Microsoft executives have repeatedly stated that the company is "all in" when it comes to cloud computing. But the software giant also continues to rake in record revenues from its traditionally delivered software products. This creates an interesting divide between its current successes and the direction in which it believes the industry is heading. Rather than choosing one over the other, Microsoft is simply creating both cloud-based and on-premises-based versions of some products. This strategy has led to some confusion, of course, but it also provides customers with choice. And because Microsoft is pushing a unique and compelling hybrid deployment model, in many cases it's also possible to mix and match local on-premises servers with hosted online services.

To date, Microsoft's most successful push in this direction has been its hosted versions of Exchange Server and SharePoint, which are currently being marketed under the Business Productivity Online Services (BPOS) umbrella, though that will change to a more expansive set of services called Office 365 starting in 2011. These services follow the hybrid model in that customers can choose between on-premises servers, hosted services, or a mixture of both. It's a proven strategy, one that is both technically excellent and popular with customers. Microsoft will employ this strategy with the next generation of Windows Small Business Server (SBS), the company's integrated solution for smaller businesses.

A Little Background

Since its inception in 1997, SBS has had a modest if compelling goal: Integrate Microsoft's most popular business server products into one package that's simpler to deploy and manage than traditional servers and that's also much less expensive. SBS has always risen to this challenge, and I've often lamented the fact that the more approachable and centralized SBS management tools weren't available elsewhere in Microsoft's products.

By the time SBS 2008 was released over two years ago, the product line was a bit out of step with the then-emerging trends in cloud computing. That version of the product offered basic integration with Office Live Small Business (also soon to be swept into Office 365), which provided a public-facing storefront or other website but little else in the way of online services integration. And in an age when many small businesses were turning to inexpensive hosted email, contacts, and calendar management from Google and elsewhere, SBS still offered a more complicated on-premises infrastructure that typically required the constant care and maintenance of a Microsoft partner.

For the next-generation version of SBS, Microsoft is splitting the product in two. At the high end of the product line is a traditional offering called Windows Small Business Server 2011 Standard, which maps very closely to previous versions and provides an obvious upgrade path for existing customers. Microsoft is also introducing a lower-end version of SBS, Windows Small Business Server 2011 Essentials, that targets the smallest of small businesses and provides very little in the way of on-premises infrastructure. With this version of SBS, customers are expected to subscribe to hosted online services such as Office 365, or even to competing Google services, in order to obtain email, contacts, calendar, and collaboration services. Essentials is all about user and computer management, centralized PC and server backup, storage, and remote access only.

Despite being based on the same underlying Windows Server 2008 R2 core, neither version of SBS has much in common with the other. They feature completely different management experiences, offer different sets of capabilities, and target different audiences. In fact, my one major

concern with this product split is whether it makes sense. Looking at the market broadly, we see companies adopting cloud services in different ways that vary, yes, by company size, but also by market type and need. New businesses, which tend to be among the smallest businesses, tend to embrace the cloud more quickly because of cost concerns. But for larger companies, cloud adoption varies mostly according to needs. Companies with regulatory, privacy, or other legal concerns about data storage tend to be more wary of the cloud and of the presumed complexity of Microsoft's hybrid solutions. Companies without such concerns are discovering that they can save a lot of money and time by moving their infrastructure to cloud services.

The SBS 2011 split divides the market like this: The smallest businesses, those with 25 or fewer employees, can pick the cloud-based Essentials solution and save a lot in up-front costs, picking the cloud infrastructure pieces they want a la carte. Of course, there are advantages to going the Microsoft route (including management integration with Essentials), but there are also cost disadvantages compared to, say, Google Apps, which is about $12 cheaper, per user per year, for small businesses compared with Microsoft's Office 365.

If your company is bigger than that, SBS 2011 Standard is your choice. Like previous SBS versions, Standard is aimed at companies with 75 or fewer employees. But it also provides on-premises email, calendar, contacts, and collaboration services, which will generally require an ongoing partner contract. So it's more expensive out of the box, and more complex. The question is whether you can save money in the long run by paying for partner management oversight instead of hosted infrastructure, as you would with Essentials.

This is a big question. For businesses already on SBS, Microsoft is further muddying the waters by supporting a migration path only to SBS 2011 Standard. So, if you want to embrace the cloud, you could find yourself on your own, or at the very least employing the services of a very specialized Microsoft partner. SBS 2011 Essentials isn't just for very small businesses. It's for small businesses that have never adopted a server infrastructure of any kind.

Consider another odd bifurcation in the product lines. With Essentials, you're free to purchase and deploy on-premises versions of Exchange or SharePoint, though of course doing so would be expensive and complex. And with Standard, you're free to decommission the included Exchange and SharePoint offerings and move to the cloud, though doing that would also be expensive—you've paid for products you're not going to use—and would work outside the centralized management experience that is a hallmark of the SBS product line. That's because while Microsoft will create management add-ons for its hosted services that are targeted at Essentials customers, it won't be doing so (for obvious reasons) for Standard customers.

SBS 2011 Standard

Windows SBS 2011 Standard (previously code-named SBS 7) is the more familiar of the two products because it's a technical follow-up to the previous generation of products. It offers Server 2008 R2 (essentially the Standard edition), Exchange 2010 SP1, SharePoint Foundation 2010, and Windows Software Update Services (WSUS) 3.1. No surprises there, although it should be noted that these are all 64-bit versions of the products.

As with previous versions of SBS, the adoption of current-generation Microsoft server products brings with it a number of new capabilities, as if by osmosis. That includes such things as the nicer Outlook Web App (OWA) in Exchange 2010 SP1, as well as the inclusion of internally hosted Office Web Apps (simple, web-based versions of Word, Excel, OneNote, and PowerPoint) with SharePoint Foundation 2010.

From a management perspective, the SBS 2011 Standard console is familiar in both look and feel. It features the now well-worn summary of the network's overall health with specific callouts for security, software updates, backups, and other alerts. User and group management is a bit less complex than going directly through the Active Directory (AD) consoles but offers a familiar level of granularity that's missing from the Elements product.

As Figure 1 shows, the SPS 2011 Standard management console is similar to that of its predecessor. Deployment is also familiar, with clean install and migration options only. Unlike Essentials, SBS 2011 Standard assumes that it will sit at the center of your network and attempt to take over DNS and DHCP duties and usurp the capabilities of the typical home or small-business router. This, too, is the way SBS has operated for years, but it's worth mentioning only because SBS 2011 Essentials doesn't do this.

"Small Business Server 2011 Standard is the traditional onsite offering, or what our customers have generally asked for," SBS Senior Product Manager Michael Leworthy told me in a recent briefing. "It's a super-simplified suite of popular server products that's easy to deploy and administer." This is a fair statement, with the understanding that it will generally require some form of Microsoft partner contract, as the small businesses that this product targets won't generally hire someone specifically for IT management. Certainly, SBS 2011 Standard deployment and management is easier than "rolling your own" with Microsoft's standalone servers. It's also considerably less expensive. It is, in other words, a great update to the existing SBS product.

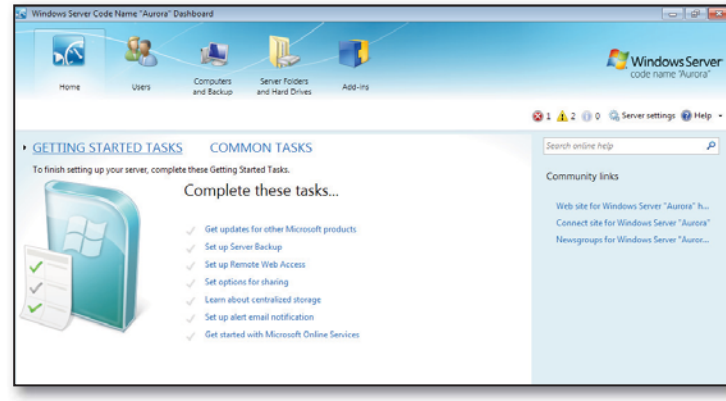

Figure 1: The SBS 2011 Standard management console

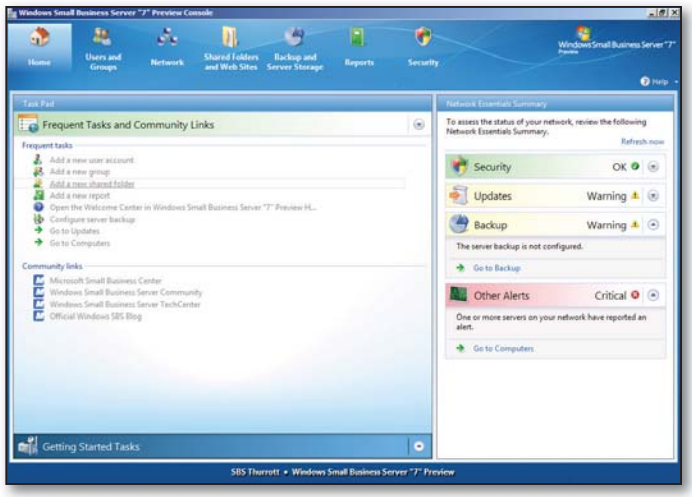

Figure 2: The SBS 2011 Essentials dashboard

It should also be noted that SBS 2011 Standard is a clear upsell for Microsoft's partners, and one with which they are already very familiar. And for businesses looking to expand, or for businesses that have specific database and line-of-business (LOB) needs, there is a new SBS 2011 Premium Add-on, which I'll describe shortly.

SBS 2011 Essentials

If self-hosting your own email and other IT infrastructure seems untenable, despite the ease of use and cost benefits of SBS 2011 Standard, take heart. For the first time, Microsoft will offer an alternative product, SBS 2011 Essentials, which is in many ways more forward-leaning and cloud-centric. Code-named Aurora, this version of SBS 2011 also runs on a Server 2008 R2 core, but it dispenses with virtually everything else found in SBS 2011 Standard. As a result, it is considerably less expensive up front, though the cost of ongoing cloud subscriptions for email and other services will have to factor into any budget.

But the biggest advantage of SBS 2011 Essentials is simplicity. Based on the same underlying infrastructure as the next Windows Home Server (WHS) product (code-named Vail), Essentials is almost too simple. I mean this in the sense that it masks some incredible complexities, which is great, but it also masks some of the more granular functionality that we've come to expect from Windows Server-based products. For example, in Aurora, there are only two types of users available from the management console: standard users and admins. Compare this with the numerous user types in SBS 2011 Standard, which again map more closely to the system's underlying user account types.

If you've worked with WHS, SBS 2011 Essentials will be familiar. However, under the covers, it's working with AD domain objects and not a simpler (but less manageable) workgroup. Indeed, the SBS 2011 Essentials and Vail management consoles are almost identical and can certainly be utilized by in-house personnel on an ongoing basis, making the long-term partner picture a bit hazy. Users can be configured for shared folder and remote web access from a simple UI. Connected computers—and the server itself—can be centrally and automatically backed up, a capability that's far more seamless and capable in this product than in SBS Standard.

SBS 2011 Essentials is as simple as it can be, with few granular options but plenty of automated capabilities that will keep the smallest businesses up and running. Figure 2 shows the SBS 2011 Essentials dashboard.

SPS 2011 Premium Add-On

In addition to the two new core SBS products, Microsoft is also delivering something called the SBS 2011 Premium Add-on. This add-on requires SBS 2011 Essentials or SBS 2011 Standard and provides licenses for both Server 2008 R2 Standard and SQL Server 2008 R2 Standard Edition for Small Business. The latter product is unique to this release and maps identically to SQL Server 2008 R2 Standard Edition from a functional perspective. The only difference is that it's licensed for use in this package with SBS only. These products must be installed together on a separate server and cannot be installed on the primary SBS machine, regardless of which version you're using.

Availability and Pricing

SBS 2011 Standard is shipping as a standalone software product and will be bundled on new server hardware starting in February 2011. It requires a server license, which costs about $1,096, plus CALs for each user, at a cost of approximately $72 each. It can be used with up to 75 users.

SBS 2011 Essentials will ship in the first half of 2011. This product was delayed somewhat because of the removal of a major feature, Drive Extender. Essentials will cost $545 for the standalone software version but doesn't require CALs for user access. It can be used with up to 25 users.

The SBS 2011 Premium Add-on will ship in December along with SBS 2011 Standard. It will cost $1,604 for the standalone software and will require Premium Add-on CALs at a cost of $92 per user. (This covers access to SQL Server only; Windows Server CALs are covered by the SBS version that you implemented.)

Final Thoughts

With this generation of products, SBS seems well poised for the future. Microsoft is doing a nice job of addressing the two most important segments of the small business market—those with in-place server infrastructure and those without—though the division between cloud-based and on-premises-based capabilities seems a bit arbitrary. Still, this is a far more aggressive play than the previous SBS version, especially if you have your eye on the cloud. What the company needs to bring it all together is cloud-based services that exceed the capabilities of the competition while meeting them on pricing. The combination of SBS 2011 Essentials with Office 365 very definitely achieves the first goal, while coming up a bit short in the pricing category.

InstantDoc ID 129147

Paul Thurrott
(thurrott@windowsitpro.com) is the senior technical analyst for *Windows IT Pro*. He writes a daily Windows news and information newsletter called *WinInfo Daily UPDATE* (www.wininformant.com).

A Conversation with Microsoft About SBS 2011

Microsoft's Kevin Kean talks about how SBS has evolved

by Paul Thurrott

With this generation of Windows Small Business Server (SBS), Microsoft is taking the venerable product line to the cloud for the first time, but it's also providing traditional SBS customers with a clear upgrade path. I spoke recently with Microsoft General Manager Kevin Kean about this change and how it will impact Microsoft's customers and partners.

Thurrott: SBS 2011 Standard has a traditional partner story, but what about the less complex SBS 2011 Essentials, which requires less oversight? Does it shut partners out or offer new opportunities?

Kean: The simplified management takes the complexity out and lets business focus on their day to day operations. Because of the platform itself, there are fewer opportunities around day to day management, sure. But we're looking at different models for partners. Certainly, there are some opportunities around reselling cloud services, and also initial deployment. But there could be a reduction in the so-called break/fix services.

We're going to offer a record of model opportunity with Office 365, where a partner gets a small business on our services and then is paid an annuity annually after that. Customers can pursue a direct relationship with Microsoft, of course. But for the partner segment, it's a resell, and one they will be paid for going forward. Partners can also customize the SBS 2011 Essentials management experience, providing more value on top of the base product.

Thurrott: You seem to be splitting the small business market very neatly with these two products. What was the thinking behind this?

Kean: SBS 2011 Essentials is targeted at that first server, to get [SMBs] up and running. Of the 36 million small businesses out there, only 9 million have servers now. A huge percentage of them see the value but don't want the cost or complexity, so they're using an old PC like a server or just getting along without. We made SBS 20011 Essentials really affordable, and it will be really affordable when bundled with hardware from an OEM or reseller.

With SBS 2011 Standard, there are 1.3 million installs of SBS 2003 and SBS 2008 out there, and we've had a large refresh cycle with our server since many of these customers have upgraded. For many, this is a good time to do a hardware refresh as well, migrate to a new server and be better positioned for the future. We've streamlined the migration process in this release, with tools that prepare the domain and investigate the environment to make sure there are no surprises. It's a straightforward proposition whether it's done in-house or by a partner.

You have to serve both markets. That said, based on the feedback we've seen, once customers start down the cloud path, they stay there. And we've seen aggressive growth of the cloud in the smallest part of the market. As these companies get bigger, they will invest more in onsite infrastructure too.

Thurrott: It seems like Microsoft is betting big on extensibility through add-ons in SBS 2011 Essentials, but not so much with SBS 2011 Standard. Is there any extensibility in SBS 2011 Standard?

Kean: SBS 2011 Standard extensibility is almost identical to that of the SBS 2008 product. But the Remote Web Access portal in SBS 2011 Standard is now the same as the one in [Essentials]. And there is some customization there, both for partners and for in-house employees.

Thurrott: There's a big change coming with the Drive Extender technologies, which were previously core to Windows Home Server and SBS 2011 Essentials (but not Standard). Can you explain that? [Drive Extender provided data duplication across two physical disks and a single pool of storage that didn't require drive letters.]

Kean: We're changing the way we handle storage in SBS 2011 Essentials. Drive Extender was a neat feature, but the implementation was off, and we discovered some application-compatibility and disk tool problems related to its ability to correct data errors on the fly. We don't want to give customers problems; we want to give them solutions. So we decided to cut out Drive Extender.

Removing Drive Extender will make file shares easy, and it's possible to accomplish most of its features otherwise. For example, you use the server's centralized backup or even RAID as an alternative to data duplication.

InstantDoc ID 129146

Microsoft Management Summit 2011

March 21–25, 2011
Mandalay Bay Hotel and Casino
Las Vegas, Nevada

You. Empowered.

Knowledge | Technology | Community

The annual Microsoft Management Summit (MMS) is the premier event of the year for deep technical content on the latest IT Management solutions from Microsoft, Partner and Industry Experts.

At this year's MMS we will share more information than ever before, covering Desktop to Datacenter to Cloud. We invite you to reserve your place at MMS 2011 where you will be empowered with in-depth knowledge and understanding of the wide range of current and soon-to-release IT infrastructure products and solutions from Microsoft and Partners.

With 150 demo-packed technical sessions and thousands of Hands-on Labs places, MMS 2011 is your best opportunity of the year to experience and evaluate the latest management technologies and connect with others in the IT management community.

Visit www.mms-2011.com today for more details and registration information.

Microsoft

FEATURE

Troubleshooting from the Wire Up for Active Directory and Beyond

Hone your troubleshooting skills to reduce backtracking, reproduce fixes, and boost your professional reputation

by Sean Deuby

As an IT pro, you've read countless articles over the years about how to troubleshoot the various moving parts of IT that make up your job responsibilities. Exchange Server, SharePoint, SQL Server, and Active Directory (AD)—each application has its own unique set of tasks and troubleshooting procedures. These applications do, however, have one trait in common: At a high level, they share a common troubleshooting methodology. Unfortunately, many IT pros don't consciously use a solid troubleshooting methodology to solve their problems and, as a result, spend more time working on issues than they need to. Understanding and consciously using solid troubleshooting principles will help you reduce the amount of time you spend backtracking in your testing, prevent you from getting lost in your troubleshooting steps, and allow you to easily reproduce fixes to your problems. I'll focus on these principles in general and later apply them to AD in particular, so you can spend less time fixing AD and more time performing useful work.

Using the Scientific Method

Troubleshooting is probably an IT pro's single most important technical skill, and yet most of us learned haphazardly how to troubleshoot. Usually, we learn from on-the-job training, perhaps from a more experienced team member or manager. It's a rare IT pro that's had professional training in structured problem solving. As a result, we have a grab bag of tools that usually works for us and a roughly structured troubleshooting process. We throw the tools at the process and hope that something we do fixes the problem.

The foundation of a solid troubleshooting skill set is to use an established methodology that guides you to the right solution. Fortunately, you don't have to invent a new methodology; the scientific method has already been invented and works for just such a task. Your first reaction is probably that you're already using the scientific method. You guess what the problem is and test your hypothesis. Did that fix the problem? If not, you go back to the first step. If you've used this method (and we all have) you've had experience with its biggest shortcoming: It's a shotgun approach. A truly effective troubleshooting methodology is more precise and detailed than this high-level approach, of course.

Cisco has created an eight-step network troubleshooting model based on the scientific method. I've expanded this model by including certain troubleshooting principles from my IT experience. The steps that follow will give you guidance in some of the more critical areas of effective troubleshooting:

1. Define the problem precisely. In this step you want to determine what the problem is exactly. Remove all vagueness and ambiguity from the problem statement. For instance, don't state, "DC1 isn't replicating correctly." This imprecise statement implies that DC1 isn't pulling updates from any upstream DC, nor are its downstream partners getting updates from it. The precisely stated problem might be much smaller: "DC1 isn't receiving updates to the configuration partition from DC2."

Precision weeds out assumptions before they can take root.

2. Gather detailed information. Ask yourself questions such as the following: What doesn't work? What does work? What changed since it was last working? If it's a client, is anyone else having this problem? If so, what do these clients have in common? Do they use the same OS build? Are they on the same subnet? Do they use the same application server? Is there anything unique about this system? If it's a server, is it behind a firewall? If so, do any other servers show the same symptoms?

3. Consider probable causes for the failure. This is the critical step in troubleshooting when you need to brainstorm and gather hypotheses, or possible reasons, the failure occurred. In a complex application such as AD, there are hundreds of potential causes for failure if you don't narrow down your choices. This is the time to apply your first troubleshooting principle, Occam's Razor, which states that, in a list of potential solutions to a problem, the simplest solution is most often the correct one. You're probably already using Occam's Razor throughout the troubleshooting process, but don't wield it carelessly; you may quickly discover that your first guess—or maybe even your second or third guess—as to the problem's root cause is wrong. Then what?

This is when you should apply a principle I call "troubleshooting from the wire up." The component you support depends on a few or many other infrastructure components. Model your troubleshooting along the lines of the seven-layer Open Systems Interconnect (OSI) model that modern networked systems are based upon, and start from the bottom, the physical network wire, up to your component. In the case of a distributed application, such as AD, the troubleshooting progression is a) the physical network, b) name-to-IP-address resolution, and c) the server OS. Check all of these before you even get to AD. When you brainstorm hypotheses, be sure to include tests to confirm that each layer is indeed working as expected. The beauty of this bottom-up method is that it works for all networked systems because that's how they're designed to work.

Now, this brute-force method may be overkill for simpler problems. If you're logged onto the server remotely, for example, you know it has power and network connectivity (at least over some ports), so you can eliminate a few steps. Just be aware that you've considered them and eliminated them. By consciously noting a step before you dismiss it, you've at least made a mental checkmark that's easier to revisit if you get stuck further in the process. Steps 2 and 3 often work in an iterative fashion; the act of gathering detailed information often reveals facts that inform new hypotheses.

4. Devise a plan to test the hypotheses. When planning your tests, as much as you may be in a hurry to bring that system back up, try to follow the next troubleshooting principle: Change only one variable at a time. If you make more than one change and the problem is fixed, you won't know whether it was the change to variable A or variable B that fixed the problem. If server A can't set up a session with server B, and you make changes to both servers' networking, you'll never know which was wrong. Sometimes the needs of the situation (or the manager) will require you to make more than one change at a time, but realize that as a result you won't be able to pinpoint the root cause.

5. Implement the plan. If at all possible, run the test with a partner or auditor. There's nothing like a second set of critical eyes to catch information the lead troubleshooter might have missed. Complex plans should always be thought through ahead of time, written down, and then closely followed. "Shooting from the hip" during the implementation might make the results invalid, leading you down a road of wrong assumptions.

6. Observe the results of the implementation. Did it fix the problem? If not, did it change the behavior of the problem at all?

7. Repeat the process from step 3, considering the next most likely hypothesis, if the plan doesn't resolve the problem. Work upward through the dependent layers of your application. If it's a SharePoint problem, is SQL Server working correctly? If it's a replication problem between two domain controllers (DCs), do they have network connectivity to each other?

8. Document the changes made to solve the problem. Once you have fixed the problem, do a post mortem to review your troubleshooting process. Some questions to ask yourself are: Did you proceed smoothly to a correct conclusion, or did you bang your head on your desk for a few hours first? If it's the latter, examine why you missed the root cause and adjust

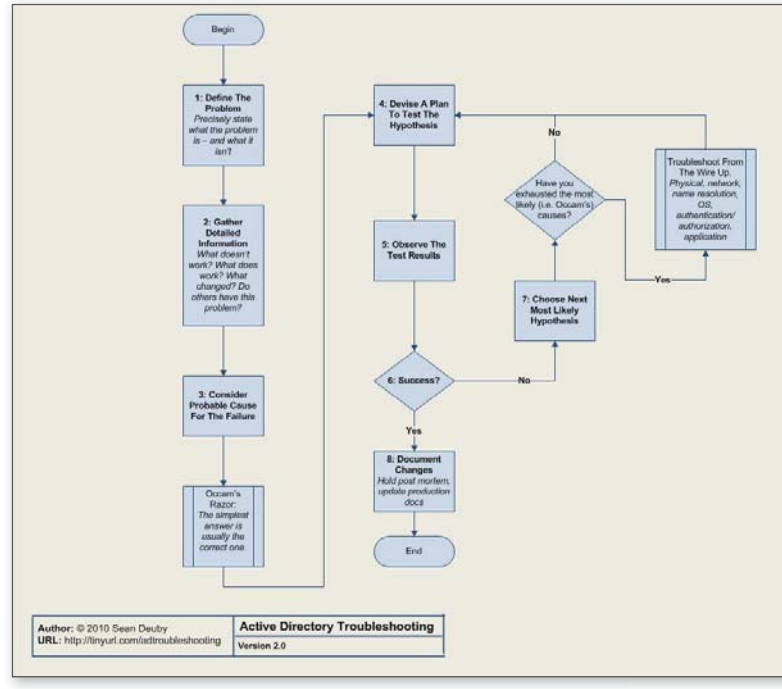

Figure 1: The eight-step troubleshooting model

your troubleshooting methodology for the next time. A post mortem doesn't need to be large and formal (unless the outage is large). What counts is that you improve your troubleshooting process so that a similar issue doesn't catch you again. Head banging is better left to rock concerts than to work surfaces!

Figure 1 refers to a flowchart for this method as part of the AD troubleshooting flowchart linked at the "AD Troubleshooting Tips & Tricks" blog at http://tinyurl.com/adtroubleshooting.

Troubleshooting AD from the Wire Up

The eight-step troubleshooting methodology applies to a wide variety of situations, both inside and outside the IT world. For the rest of this article, let's focus on troubleshooting AD and the tools you should use. AD is a complicated application that relies on other complicated components. What dependent components need to be in good shape for AD to work right?

Figure 2 shows the architectural layers most important to AD functionality: physical, network, name resolution, OS and authentication, and finally AD itself. The physical layer seems obvious—nothing works without power. Or if the network cable isn't connected, the packets can't flow—but I've lost track of the number of times operations was trying to troubleshoot an unavailable DC only to discover a site was shut down due to a national holiday, and the site operations forgot to tell central operations. This layer also encompasses hardware failures on the DC itself.

The network layer is where you should check for IP address or subnet configuration errors, WAN or LAN failures, and firewall changes that block ports used by AD. Your best tools for this layer are Ping, Ipconfig, and Tracert. This last possibility has proven to be a common (and frustrating) root cause, perhaps because the solution isn't technology driven. To correct this, what's often needed is better communication between the network and directory teams.

Besides AD itself, name resolution is the layer where you should have the strongest troubleshooting skills because that's where most of the non-AD issues are found.

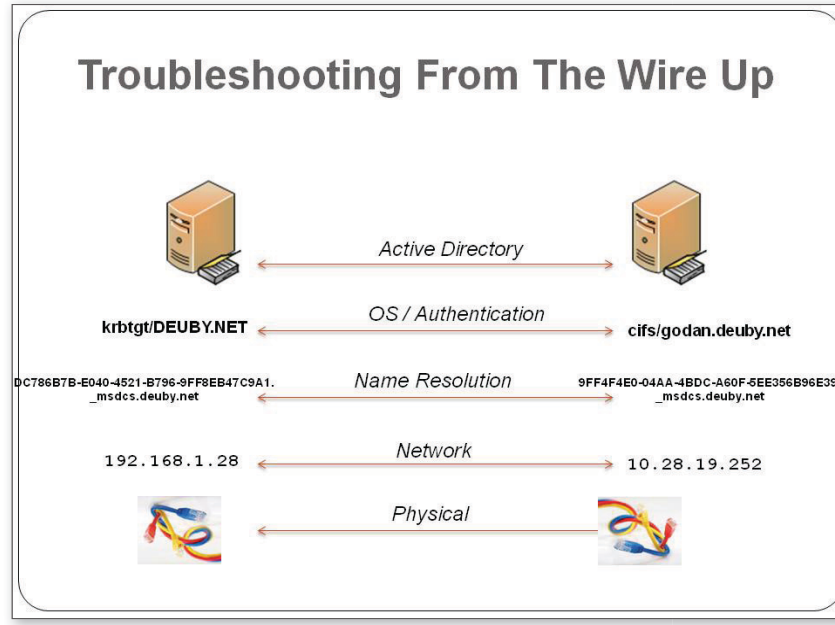

Figure 2: Architectural layers vital to AD

DNS performs the hostname to IP address resolution for AD, but it also performs a variety of other important functions, such as service location and DSA ID (a unique identifier used for replication) translation. The tools for troubleshooting issues related to name resolution are Nslookup, Nltest, Dcdiag, and Ipconfig.

Next is the OS. Dedicated DCs are less complicated to troubleshoot at the OS level than other application servers because they typically have only the base OS, Active Directory Domain Services, and DNS. In Windows 2003 Server, you can use Netsh Diag GUI and Netdiag. In Windows Server 2008, you can use Server Manager and the Performance and Reliability Monitor. Server 2008 R2 adds the Best Practices Analyzer. All versions, of course, have the event log. Finally, MPSReports, available from the Microsoft Download Center, is a configuration-gathering tool used by Microsoft Support Services to diagnose your system. You can use it, too, and view the results of the report using the MPSReports Viewer (also from the Microsoft Download Center).

Within the OS layer is authentication (i.e., Kerberos), which I've broken out because it's a far more common error source than the general OS. Kerberos relies on close time synchronization and both ticket-granting tickets and session tickets to successfully authenticate identities to resources in the domain. The tools in this area are the event log, Kerbtray, and Klist (from the Microsoft Download Center). Sometimes the problem isn't related to any failures in the authentication mechanism itself but to human "assistance." I've had DCs begin failing authentication because certain Latin American countries legislated unique daylight saving time (DST) changes. The DC in that country didn't have the hotfix that recognized the change, so operations changed the time manually. This time skew caused the DC's Kerberos session tickets to fail, and the DC began throwing errors.

Resolving issues in AD requires all these skills to help isolate the various moving parts of this distributed application. As an IT pro you're called upon to troubleshoot a wide variety of systems, from an AD implementation that supports thousands of users, to your parent's home computer, to the kitchen light switch. If you build yourself a strong foundation of structured, logical troubleshooting skills, you can repair all of these situations and boost your professional reputation at the same time.

InstantDoc ID 128973

Sean Deuby
(sean@windowsitpro.com) is technical director for *Windows IT Pro* and *SQL Server Magazine*, and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.

FEATURE

A Concentrated Guide to PowerShell Functions

Modularize a set of commands

by Don Jones

Windows PowerShell 2.0 offers several ways to modularize a set of commands. Solutions range from easy to complex, depending on your needs. However, all the solutions have some common rules and traits.

First, an individual function should ideally be focused on a single task. A function should either produce no output (as might be the case if it's taking some action), or it should produce a single kind of output (as might be the case if it's retrieving some information). Using cmdlet-style verb-noun naming for your functions is one way to help keep them single-tasked.

In addition, functions should create output only by using the Write-Output command. If you need a function to also log something to a file (e.g., errors), use Out-File. You can use Write-Verbose, Write-Debug, Write-Warning, and Write-Error to output such information (verbose progress, debug information, warnings, and errors, respectively). Avoid using Write-Host because its output can't be used as flexibly as the other Write cmdlets' output can.

Finally, functions' input should come entirely through parameters. A function should never refer to any variables created outside itself.

To illustrate the power of PowerShell functions, let's consider a situation in which we want to use Windows Management Instrumentation (WMI) to retrieve OS information and BIOS serial numbers from one or more remote computers. Thus, from the Win32_OperatingSystem class, we want the Caption, ServicePackMajorVersion, and BuildNumber properties; from the Win32_BIOS class, we want the SerialNumber property. We also want the output to contain the computer names.

Start with a Command

I prefer to work directly from the command line, rather than in a script file, to get the main functional portion of a task out of the way first. The following two commands obtain the information we're looking for:

```
Get-WmiObject -class Win32_OperatingSystem -computername Server-R2 |
  Select __SERVER,ServicePackMajorVersion,BuildNumber,Caption
Get-WmiObject -class Win32_BIOS -computername Server-R2 | Select SerialNumber
```

You can add a comma-separated value (CSV) list of computer names to -computername to target additional computers.

A Single Kind of Output

A problem with running two commands is that you get two separate sets of output, which isn't what we want the eventual function to do. So, we need a way to combine the desired information into a single output table.

Listing 1: Script to Store WMI Information in Two Variables

```
$os = Get-WmiObject -class Win32_OperatingSystem -computername Server-R2
$bios = Get-WmiObject -class Win32_BIOS -computername Server-R2
$obj = New-Object PSObject
$obj | Add-Member NoteProperty ComputerName ($os.__SERVER)
$obj | Add-Member NoteProperty OSVersion ($os.Caption)
$obj | Add-Member NoteProperty OSBuild ($os.BuildNumber)
$obj | Add-Member NoteProperty SPVersion ($os.ServicePackMajorVersion)
$obj | Add-Member NoteProperty BIOSSerial ($bios.SerialNumber)
```

It would be nice to call the __SERVER column ComputerName rather than just __SERVER. Using the name __SERVER is handy because __SERVER is a WMI system property, and the __SERVER command always returns a computer's real name (regardless of what nickname or IP address you specified to reach the computer)—but __SERVER is ugly.

As is usually the case in PowerShell, numerous methods exist for accomplishing this task. One solution is to create a new, blank object and add just the properties you want to it from Win32_OperatingSystem and Win32_BIOS. This approach is kind of a script-based solution, but the logic is pretty clear. The script in Listing 1 lets you store the WMI information in two variables (one for each WMI class), then specify the bits to tack on to a new, blank object.

An entirely different approach is to use the Select-Object command. This command creates a new custom object with whatever properties you want. You can use a special trick called a hashtable or dictionary to add a property from Win32_BIOS. This solution is a one-line command rather than a full script, but the syntax—especially the punctuation—is a bit harder to follow. The following command renames the __SERVER property to ComputerName but leaves the other property names as they are.

```
Get-WmiObject -class Win32_
  OperatingSystem -computername
  Server-R2 | Select @{n=
  'ComputerName';e={$_
  .__SERVER}},Caption,BuildNumber,
  ServicePackMajorVersion,
  @{n='BIOSSerial';e={Get-WmiObject
  -class Win32_BIOS -computername
  $_ | Select -expand SerialNumber}}
```

I know it seems crazy, but it works!

For the purposes of this article, let's use the solution presented in Listing 1. I prefer this solution because the syntax is easier to follow.

Simple and Parameterized Functions

The simplest type of function simply wraps your commands into a function construct, as the script in Listing 2 does. This script uses the Write-Output command to actually output the custom object to the pipeline. The problem with this function is that someone must still open it and edit it to change the computer name, which is a hassle.

Adding a parameter to the function lets someone else run it without having to edit it. The script in Listing 3 declares the parameter as a string, as well as provides a default value of 'localhost' in case someone forgets to provide the parameter. The parameter gets dropped in place of the hard-coded computer name on the two Get-WmiObject commands.

Several options exist for running this script. You can use a positional parameter, like so:

```
Get-OSInventory Server-R2
```

or a named parameter, like so:

```
Get-OSInventory -computername Server-R2
```

Either way, pulling the function into the shell to use it as a command is a bit tricky.

Converting a Function into a Command

You have two simple options for making the new function visible in the shell (plus a third option that I discuss later in the article). The first option is to copy the function into a profile script that executes automatically every time you open a new shell window. For information about using PowerShell profiles, run the command

```
help about_profiles
```

The profile scripts you'll use don't exist by default; the Help file tells you what to name them and where to put them. Just copy the function straight into a profile script to run it.

Another option is to dot source the script; for example:

```
. c:\Scripts\Utilities.ps1
```

This command loads the script's contents into the shell's scope, making whatever's in the script available globally throughout the shell—until you close the shell, of course. A disadvantage of this option is that you must rerun the command every time you open

Listing 2: Simple Function that Wraps Commands into a Function Construct

```
Function Get-OSInventory {
 $os = Get-WmiObject -class Win32_OperatingSystem -computername Server-R2
 $bios = Get-WmiObject -class Win32_BIOS -computername Server-R2
 $obj = New-Object PSObject
 $obj | Add-Member NoteProperty ComputerName ($os.__SERVER)
 $obj | Add-Member NoteProperty OSVersion ($os.Caption)
 $obj | Add-Member NoteProperty OSBuild ($os.BuildNumber)
 $obj | Add-Member NoteProperty SPVersion ($os.ServicePackMajorVersion)
 $obj | Add-Member NoteProperty BIOSSerial ($bios.SerialNumber)
 Write-Output $obj
}
```

Listing 3: Utilities.ps1

```
Function Get-OSInventory {
 Param([string]$computername = 'locahost')
 $os = Get-WmiObject -class Win32_OperatingSystem -computername $computername
 $bios = Get-WmiObject -class Win32_BIOS -computername $computername
 $obj = New-Object PSObject
 $obj | Add-Member NoteProperty ComputerName ($os.__SERVER)
 $obj | Add-Member NoteProperty OSVersion ($os.Caption)
 $obj | Add-Member NoteProperty OSBuild ($os.BuildNumber)
 $obj | Add-Member NoteProperty SPVersion ($os.ServicePackMajorVersion)
 $obj | Add-Member NoteProperty BIOSSerial ($bios.SerialNumber)
 Write-Output $obj
}
```

Listing 4: Pipeline Function

```
Function Get-OSInventory {
 BEGIN {}
 PROCESS {
 $computername = $_
  $os = Get-WmiObject -class Win32_OperatingSystem -computername $computername
  $bios = Get-WmiObject -class Win32_BIOS -computername $computername
  $obj = New-Object PSObject
  $obj | Add-Member NoteProperty ComputerName ($os.__SERVER)
  $obj | Add-Member NoteProperty OSVersion ($os.Caption)
  $obj | Add-Member NoteProperty OSBuild ($os.BuildNumber)
  $obj | Add-Member NoteProperty SPVersion ($os.ServicePackMajorVersion)
  $obj | Add-Member NoteProperty BIOSSerial ($bios.SerialNumber)
  Write-Output $obj
 }
 END {}
}
```

the shell (unless you put the dot-sourcing command into your profile).

Accepting Pipeline Input

The ability to send pipeline input to the function would be useful, but the function can currently accept only a positional or named parameter. Listing 4 includes a special kind of function, called a pipeline function or filtering function, that sends pipeline input to a function.

Assuming you have a text file full of computer names, you'd run the function with one name per line, like so:

```
Get-Content names.txt | Get-OSInventory
```

When the names get piped into the function, PowerShell executes the BEGIN block first. I don't have any commands in that block, but I could use the block to set up a database connection or something else I wanted to reuse throughout the function.

The shell executes the PROCESS block once for each name piped in. It puts each computer name from the text file into the $_ placeholder—which I recommend copying into the $computername variable, for better readability.

Finally, when all the names have made it through the PROCESS block, the END block executes. Again, I didn't need to use the END block, but I could have closed a database connection or performed other cleanup tasks before the function finally completed. You can omit the BEGIN and END blocks if you don't want to use them, but I like to include them to maintain consistency with my other functions.

An advantage of using a function this way is that the approach works with any technique to obtain computer names because the act of getting the names is external to the function itself. For example:

```
Get-ADComputer -filter * -searchbase
  'ou=West,dc=company,dc=com' } |
  Select -expand Name |
  Get-OSInventory
```

This command uses the Windows Server 2008 R2 Active Directory (AD) module's Get-ADComputer command to obtain computers from the West organizational unit (OU) of the company.com domain. The Select-Object command retrieves just the contents of those computers' Name properties, piping the computer names into the function. The benefit of the function not worrying about where the computer names came from is that you can reuse the function, without alteration, with a variety of computer name sources. Keeping a function single-tasked helps ensure that you can reuse the function in multiple places.

> PowerShell 2.0 provides a solution called an advanced function, which is also called a script cmdlet because it works similarly to a real cmdlet.

Having a function output custom objects to a pipeline makes the function more flexible, too. For example, suppose you want the output in a CSV file:

```
Get-ADComputer -filter * -searchbase
  'ou=West,dc=company,dc=com' } |
  Select -expand Name |
  Get-OSInventory | Export-CSV
  output.csv
```

Or perhaps you want to filter the output so that you only get Windows XP computers running something other than SP3:

```
Get-ADComputer -filter * -searchbase
  'ou=West,dc=company,dc=com' } |
  Select -expand Name | Get-
  OSInventory | Where { $_.OSBuild -eq
  2600 -and $_.SPVersion -ne 3 }
```

Keeping a function single-tasked and outputting the objects lets the function work with a variety of other PowerShell commands.

Now you can accept pipeline input. However, note that because you deleted the parameter block, you also removed the ability to use the -computername parameter. (Big sigh.) What you really need is a function that could work either way—just like a cmdlet.

Advanced Functions: PowerShell's Script Cmdlets

It would be nice to have a function that could support both pipeline input and the use of positional and named parameters. PowerShell 2.0 provides a solution called an advanced function, which is also called a script cmdlet because it works similarly to a real cmdlet. Advanced functions are definitely *advanced* in nature.

A significant problem lies in the fact that if you pipe input into the function, the function will execute the PROCESS script block one time for each input item. Each item is then placed, one at a time, into the parameter.

Using a parameter causes the PROCESS script block to execute only once. However, the parameter contains all the input items, which you must manually enumerate through.

I like to create advanced functions as a kind of shell that's designed to handle both of these scenarios but that doesn't do any real work. A better solution is to move the actual work into a support function that's designed to work with one computer at a time. The advanced

Listing 5: Advanced Function to Query OS Information

```
Function OSInventoryHelper {
 Param([string]$computername)
 $os = Get-WmiObject -class Win32_OperatingSystem -computername $computername
 $bios = Get-WmiObject -class Win32_BIOS -computername $computername
 $obj = New-Object PSObject
 $obj | Add-Member NoteProperty ComputerName ($os.__SERVER)
 $obj | Add-Member NoteProperty OSVersion ($os.Caption)
 $obj | Add-Member NoteProperty OSBuild ($os.BuildNumber)
 $obj | Add-Member NoteProperty SPVersion ($os.ServicePackMajorVersion)
 $obj | Add-Member NoteProperty BIOSSerial ($bios.SerialNumber)
 Write-Output $obj
}
Function Get-OSInventory {
 [CmdletBinding()]
 Param(
  [Parameter(Mandtory=$True,ValueFromPipeline=$true)]
  [String[]]$ComputerName
 )
 BEGIN {
  $inputWasFromPipeline = -not $PSBoundParameters.ContainsKey('computername')
 }
 PROCESS {
  If ($inputWasFromPipeline) {
   OSInventoryHelper $computername
  } else {
   foreach ($computer in $computername) {
    OSInventoryHelper $computer
   }
  }
 }
}
```

function ensures that the input is broken down into one computer at a time if necessary. Listing 5 contains the code for the advanced function.

A complete discussion of advanced functions is beyond the scope of this article. For more information about these fun constructs, run the command

```
help *advanced*
```

You'll notice that the real work in the script in Listing 5 is done in the OSInventoryHelper function. You can hide the OSInventoryHelper function from users, forcing them to use Get-OSInventory directly. (I explain how to do so in the next section.)

The advanced function in Listing 5 can be used in multiple ways. For example:

```
Get-Content names.txt | Get-OSInventory
  | Export-CSV output.csv
Get-OSInventory -computername
  server-r2,server57,dc001 |
  Format-Table
Get-OSInventory localhost
Get-OSInventory
```

The last command actually prompts you for one or more computer names; press Enter on a blank line when you're done entering computer names.

Distributing the Result

Now that we have an advanced function, we need to distribute it in such a way that other administrators can easily load it into the shell without having to use a profile script or dot sourcing. In addition, it would be nice to hide the OSInventoryHelper function.

We can accomplish both tasks by saving OSInventoryHelper and the actual Get-OSInventory functions into a file called Utilities.psm1 that includes the following additional line of code

```
Export-ModuleMember -function
  Get-OSInventory
```

to ensure that *only* Get-OSInventory is visible to other administrators. (OSInventoryHelper is visible if you open the file, but it's hidden when you load the script into the shell.) The Utilities.psm1 script should be saved in the \My Documents\WindowsPowerShell\Modules\Utilities folder because the shell automatically searches this path for new modules.

To load Get-OSInventory into the shell as a command, run

```
Import-Module Utilities
```

The road to this advanced function within a script module was long and full of complex syntax, but the end result is beautiful.

InstantDoc ID 128912

Learning Path

For more information about PowerShell functions:

"Create Your Own PowerShell Functions," InstantDoc ID 101610

"Customizing Windows PowerShell's Internal Functions," InstantDoc ID 98314

"Q. How do I make a PowerShell function support ShouldProcess?" InstantDoc ID 125333

"Q. How can I make a PowerShell script or function have mandatory parameters?" InstantDoc ID 125336

Moving from command lines to scripts:

"Editing and Debugging Scripts with PowerShell 2.0's Integrated Scripting Environment," InstantDoc ID 104713

"Q. Is there such a thing as a 'script cmdlet' in PowerShell?" InstantDoc ID 125724

"Error Trapping and Handling in PowerShell," InstantDoc ID 125327

"Export PowerShell History for Instant Scripts," InstantDoc ID 125545

"Debugging in Windows PowerShell," InstantDoc ID 125694

"Protect Your PowerShell Scripts," InstantDoc ID 102831

For information about PowerShell profiles:

"What You Need to Know to Start Using PowerShell's Personal Profile Scripts," InstantDoc ID 97669

"Save Your PowerShell Code in Profile and Script Files," InstantDoc ID 101718

You can find many more tips from Don Jones in his PowerShell with a Purpose blog at www.windowsitpro.com/blogs/PowerShellwithaPurpose.aspx.

Don Jones
(powershell@concentratedtech.com) is the author of more than 35 books, and is a speaker at technology conferences such as Microsoft TechEd and Windows Connections. He's a multiple-year recipient of Microsoft's MVP and is technical guide for PowerShell at www.windowsitpro.com.

Virtual Desktop Infrastructure, Part 1: Everything *Except* VDI

Desktop virtualization is the first step

by John Savill

Virtual Desktop Infrastructure (VDI) is a hot topic these days. VDI is a form of desktop virtualization in which the user's entire desktop experience is hosted in the data center. The user remotely connects to the desktop from some type of client device; none of the user's desktop, applications, or data actually resides on the local device.

I have several upcoming articles in which I discuss what VDI is, where it fits best, the technologies involved in VDI, creating the right infrastructure, and how Citrix's XenDesktop enhances Microsoft's approach to VDI. In this first article, however, I focus on some technologies that at first glance seem to have nothing to do with VDI. These desktop virtualization solutions aren't actually part of VDI but are critical components of a successful VDI architecture.

Desktop Environment

The most important aspect of a user's computer experience is the applications, which perform functions and manipulate both local data and data on servers. The OS is the primary tool for running applications and managing data. Users often customize their OS environment with special backgrounds, screen savers, shortcuts, and favorites. Although these changes might seem disposable to an administrator, users might spend hours trying to find a lost application or data, or recreate their shortcuts. Thus, maintaining the user's environment is important.

Three main areas comprise the desktop environment: user data and settings, applications, and the OS. The underlying hardware that the OS is installed on is also important.

In most desktop environments, all these components are intertwined rigidly with one another. The OS is locally installed on the user's desktop computer; the applications are installed directly onto the OS, making changes to the file system, registry, and other OS components; and the user data and settings are stored on the local file system. These layers are often depicted as Figure 1 shows because we typically install the applications onto the OS, the user customizes the OS and applications, and the user has data—but in reality, the applications and user settings and data are actually bound onto the OS layer.

This tight, localized coupling introduces several problems:

- Having data solely locally on a client machine introduces the risk of data loss because of hardware loss, hardware failure, volume corruption, or accidental deletion. Data is also inaccessible if the user utilizes alternative hardware. The same problem applies to user settings and configuration.
- A failure of the OS or desktop hardware means complex procedures to recover information and settings. Installed application listings must be obtained before the hardware can be replaced or the OS reinstalled; then all applications must also be reinstalled.

Figure 1: Typical desktop layers

- An application failure results in complex troubleshooting and uninstallation processes because of changes that must be made in multiple places on the OS.
- Deploying applications and application updates can be very complex and time consuming.

The solution to all these problems is to virtualize each of the layers, making them independent and abstracted from one another. This solution offers a flexible environment that's easy to deploy, easy to maintain, and easy to enable a consistent "anywhere access" experience.

The desire to separate elements of a system to make it more flexible is common: Many products advertise an easy-to-switch modular design. We've all bought those all-in-one TV/DVD/video combinations only to curse when the video part breaks and we're left without TV or DVD while the video component is repaired. But with separate TV, DVD, and video components, it's easy to swap out one component without losing functionality of the others. That's what we want for the user's desktop experience—for the OS, applications, user settings, and data to be separate blocks that are pulled together and assembled depending on the logon environment. This environment could be a local desktop, a VDI OS, or a Terminal Services/Remote Desktop Services session. Because the components are separate blocks, there's no delay waiting for installation or configuration. The components are just layered on top of each other to provide the full desktop experience.

In this article I discuss the technologies that let us separate these typically heavily intertwined layers to allow on-the-fly assembly that produces a complete desktop environment no matter where a user logs on. Figure 2 shows a sample environment that includes various solutions.

User Data and Settings

The goal is to be able to extract a user's settings and data from the desktop environment so that the settings and data are protected and available no matter where the user connects—for example, the user's own desktop, someone else's desktop, or a remote desktop session to a corporate presentation virtualization solution such as Terminal Services (Remote Desktop Services), Citrix's XenDesktop, or a VDI-hosted remote client OS. We've had technologies to handle the abstraction of user settings and data for a long time in Windows, but with Windows 7 the technologies are tuned and enhanced to a functionality and experience level that doesn't negatively affect the user experience but instead improves the settings and data availability.

Let's consider user settings first. Each user has a profile on his or her machine, under the C:\Users folder for Windows Vista and later. This profile consists of several files and folders, including the ntuser .dat file that contains all the user-specific registry information. Although this file is small, it constitutes the bulk of a user's customization. The profile also contains Internet Favorites, documents, searches, and other types of information. I cover data items later in the article.

To achieve a consistent user experience, a user's profile must be available no matter where the user connects. Thus the profile must be stored on the network. This capability, called roaming profiles, has been possible for a long time in Windows.

In the past, organizations were reluctant to use roaming profiles because of how they worked. A user's profile was simply uploaded to the network during logoff, which caused a long delay in logging off. Windows 7 introduced background synchronization of roaming profiles. This feature is disabled by default, but you can use Group Policy to enable it. Background synchronization syncs the user's profile at specific times or at a certain periodic time interval. Periodically synchronizing data means there's less data to synchronize

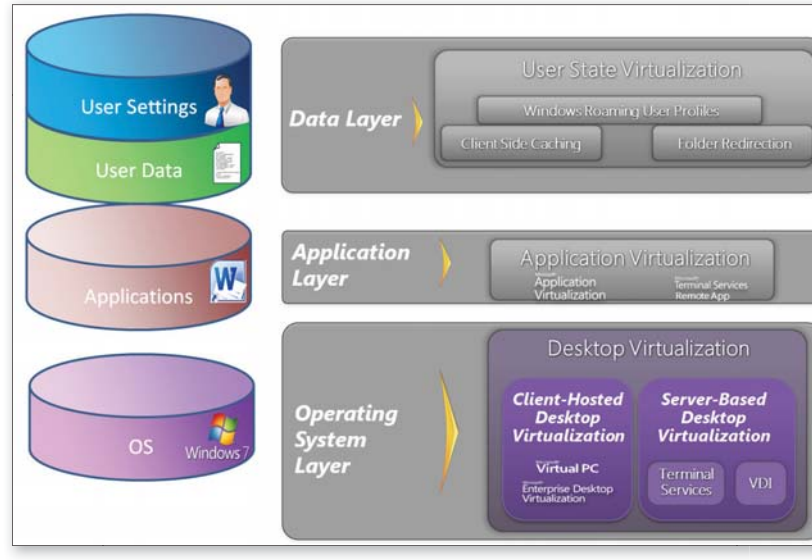

Figure 2: Desktop virtualization components

when it's time to log off, which results in a far smoother end-user experience.

Another reason that roaming profiles can be problematic is that before Windows 7, user data was left as part of the profile, which resulted in a huge profile. In reality, roaming profiles were never designed to handle the replication of user data. Even with the advancements in profiles in Windows 7, the user data must still be stripped from profiles—which leads us to the topic of folder redirection.

Several data storage locations are available to users, such as the Documents and Pictures folders, which by default are subfolders of the user's profile. If roaming profiles are used, these folders and the data they contain are replicated as part of the profile replication process, which as we already established isn't optimal. Folder redirection is an alternative technology that lets us configure the standard folders to point to a specific location on the network. For example, when a user accesses the Documents folder, he or she is actually accessing a network location—although this fact is entirely transparent to the user. You use Group Policy to configure folder redirection.

One of the biggest changes from Windows XP to Vista is the restructuring of the user profile namespace to allow greater separation of the various types of data. In addition, the number of distinct profile folders that can be redirected increased from 5 distinct storage areas in XP to 13 in Vista, which includes separation of the Documents, Pictures, Videos, and Music folders. In XP, Pictures, Music, and Videos are all subfolders under Documents and therefore follow that folder's redirection configuration. In Vista, you can opt to *not* redirect the Music and Videos folders if you choose. Folder redirection now lets you redirect all the user data you want, which makes the user profile very small and lets roaming profiles easily handle the remaining data synchronization (i.e., ntuser.dat and some other minor data files).

The Offline Files feature (also known as Client Side Caching) lets you store a local copy of the user's data from the network on certain configured machines that still require access to data even when disconnected from the network (e.g., laptops). Offline Files synchronizes the data changes at a delta level when network connectivity is restored. This delta-based replication means that only changes to files are replicated, instead of the entire file.

Applications

When we build a machine, the first thing we do after installing the OS and its updates is install the applications, which can include Microsoft Office, line of business (LOB) applications, security services, and other types of software. Installing software typically takes a significant amount of time because of the setup routines and configurations required to update file systems, add registry values, and register resources. Typically, only organization-wide applications are installed during OS installation. Applications that are user or department specific might be installed at first logon, which adds further delays.

In addition to time delays, the following installation problems can occur:

- Application-to-application compatibility—Because of how applications modify the OS, they often cause incompatibilities with other applications or even with different versions of the same application. Thus significant regression testing is necessary before putting a new or updated application into production, and certain application combinations might be prohibited.
- OS bloat—As each piece of software is installed, extra services are added that use resources but that don't always provide value. In addition, the registry increases in size, which uses memory and slows down the system. Even when applications are later uninstalled, pieces are often left behind on a system.
- Application updates—The processes for updating applications can vary widely, which requires significant planning and infrastructure for deployment.

These problems are all related to the fact that applications are installed on the OS. Imagine users logging on to different computers, remote sessions, and VDI environments—and all needing different applications. Installing every application that every user might ever need on every OS environment simply isn't practical and would result in a hugely bloated OS that would be a nightmare to maintain because every application update would have to be applied to every OS instance. Using traditional software installation methods as users log on to different environments isn't feasible because of the time expense, not to mention the additional problem of uninstalling applications on logoff.

One solution is traditional presentation virtualization, in which applications are installed on terminal servers, then executed on the terminal server while the application's window displays on the user's local desktop. This solution requires a significant server infrastructure to host the application execution and prohibit offline application execution. In addition, we still have the problem of all the applications needing to be installed on several terminal servers. However, this solution might work for certain applications—for example, an application that requires access to large amounts of data that's housed in the data center. When such an application is run in the data center via presentation virtualization, the network traffic associated with the data access is restricted to the data center network, which is typically very fast. Running the same application locally on a user's desktop sends all the data over the network, which uses a lot of bandwidth and slows the application execution.

The other major application solution is application virtualization. The big difference between application virtualization and presentation virtualization is that in application virtualization, the applications actually execute on the local OS instance. And even more importantly, the applications execute without needing to be installed on the local OS thanks to a per-application virtual environment that lets applications execute without changing the local OS.

Remember that when an application is installed, changes are made to the file system (e.g., the application's executables and DLLs are placed in C:\Program Files, changes are made to the registry, services are installed). Application virtualization works by capturing all these system changes during a process known as sequencing (Microsoft App-V's term), which involves converting an application's installation routine to a binary stream that can be used with application virtualization.

With sequencing, system changes are captured during application installation. These changes are saved in virtual layers, such as the file system layer, registry layer, services layer, fonts, OLE, and configuration, which can then be loaded into the virtual environment when the application launches. The application thinks its files, registry, and services are all on the local OS but in reality the application just sees the virtual layers that the application virtualization technology facilitates. Nothing is actually written to the underlying OS.

Figure 3 shows the application virtualization layers. Note that although virtualized applications can't write to OS resources beyond such items as user configuration and data, they can read information from the host OS.

App-V is Microsoft's application virtualization solution. By default, App-V works as a streaming technology. The first time a virtualized application is launched on an OS, the App-V client communicates with an App-V streaming server that sends to the client the part of the application's binary stream that's necessary to initially launch the application. This portion of the binary stream is known as Feature Block 1 and is typically about 10 to 20 percent of the total binary stream. It's sent to the client very quickly—in Office 2010, about a 3-second delay occurs between the user clicking an application icon and the application window launching.

The App-V sequencing process determines what needs to be placed in Feature Block 1 by actually launching the application during the sequencing and monitoring the parts of the stream needed. The necessary items go to Feature Block 1 and the rest goes to Feature Block 2. After the application launches, Feature Block 2 is sent to the client in the background. This binary stream is cached on the local client OS, so if the application is launched again the stream doesn't have to be sent over the network again.

Although I said application virtualization makes no changes to the local OS, it obviously caches this stream. However, the stream is cached into one file in the All Users profile and this single cache file is shared by all users and applications. Applications don't write anything anywhere else in the file system or make any configuration or registry changes. This cache also means that virtualized applications are available even when the machine is offline. In a VDI environment we can actually configure this cache to be placed on a file share that's common to all the VDI client virtual machines. This approach eliminates the need for each client to have its own App-V cache, which saves disk space and expedites the initial application launch.

App-V's default distribution method is streaming. However, App-V also supports the creation of an MSI file that contains the complete stream for other deployment technologies, such as Group Policy and Microsoft System Center Configuration Manager (SCCM) 2007 R2. We can even leverage traditional file shares and Microsoft IIS for App-V stream distribution. A lot of options are available to suit different organizational needs and infrastructures.

Note that basic cut and paste, object linking, and embedding still work between virtualized applications, but for deeper integration we can create dependencies between virtualized applications. This dynamic suite composition lets separate virtualized applications see each other in a controlled manner.

Application virtualization cures several deployment problems. For example:

- Applications can be rapidly deployed on a per-user, as-needed basis, effectively in real time, which lets administrators slot applications into the desktop environment as necessary.
- Application-to-application compatibly issues are solved because separate virtualized applications no longer see one another. Applications have their own unique virtual file systems, registries, etc. This isolation also removes most of the regression testing needed when introducing new or updated applications.
- Rolling out updates is a simple process. The sequenced application is updated only once, and App-V rolls out the changes to the stream to all clients, without any user action.
- The OS doesn't suffer from bloat because applications aren't actually installed on the OS.

Operating System

Organizations that are moving to Windows 7 might face application compatibility problems. If a newer version of an application isn't available that's compatible with Windows 7, or a similar product doesn't exist, OS virtualization is an option. Several ways exist to virtualize the client OS.

Windows 7 introduced Windows XP Mode, which allows applications that won't run on Windows 7 to execute on a local XP virtual machine (VM). The application window seamlessly displays on the user's main Windows 7 desktop, totally transparently to the user. In such a case, we're still running a separate legacy OS on the user's Windows 7 desktop that needs to be managed. Microsoft Enterprise Desktop Virtualization (MED-V), which is part of the Microsoft Desktop Optimization Pack (MDOP), simplifies this process by providing a centralized method to not only distribute and update the XP VM but also manage shortcuts on the desktop and URL redirections to Internet Explorer (IE) 6.0. This approach provides the added benefit of solving Windows 7's IE compatibility problem.

Another approach to client virtualization is to virtualize the user's entire desktop

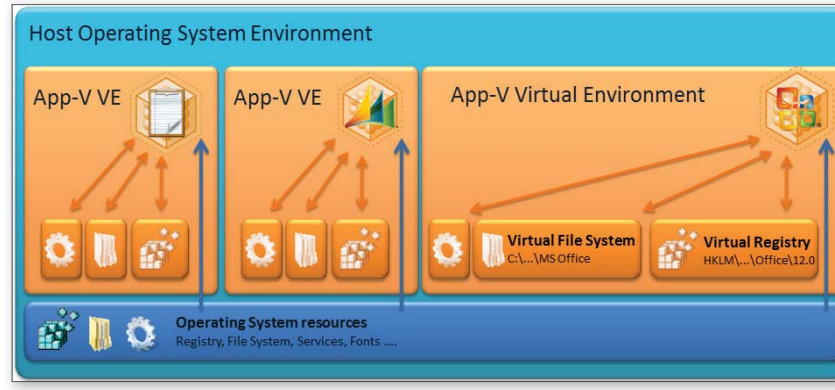

Figure 3: Application virtualization layers

OS in the data center, giving the user a local client to enable remote connectivity to the data center–hosted client OS. This local client could be a thin device, a legacy PC running Windows Fundamentals for Legacy PCs, or any other type of device or OS that supports the RDP protocol (in the case of a pure Microsoft solution). A benefit of this approach is that because the user's entire desktop is housed in the data center, sensitive data never actually leaves the data center. In addition, the desktop is available no matter where the user connects, including from a personal machine at home. This solution is great for disaster planning.

Putting It All Together

When a user logs on to a new OS instance for the first time, whether it's a local fat desktop, a session on a terminal service, or a VDI-hosted client OS, the following steps occur:

1. The user logs on to the OS using his or her Active Directory (AD) account.
2. After user authentication occurs, the parts of the profile that weren't abstracted through folder redirection are pulled down; this minimal amount of information downloads very quickly. All the customizations are now present in the user's session, in addition to the folder redirection settings. Thus all the user's data, favorites, and so forth are present.
3. The App-V client communicates with the App-V management server to determine the applications that apply to the logged on user and subsequently populates shortcuts on the desktop and Start menu, in addition to configuring the relevant file type associations.
4. The user now has a fully functional desktop and can launch applications and access data with no delays.

Desktop virtualization and VDI aren't the same. VDI leverages desktop virtualization technologies to provide a data center–hosted client OS. Although VDI might be the best solution for certain users, every desktop environment can benefit from some form of desktop virtualization—whether it's application virtualization, user state virtualization, or OS virtualization. In planning your desktop architecture, especially as part of a Windows 7 rollout, make sure desktop virtualization is considered as part of that architecture. The additional upfront work yields huge long-term benefits. Desktop virtualization not only provides you with a more agile environment but also saves infrastructure and management costs.

InstantDoc ID 129007

John Savill
(john@savilltech.com) is a Windows technical specialist, an 11-time MVP, and an MCITP: Enterprise Administrator for Windows Server 2008. He's a contributing editor for *Windows IT Pro,* and his latest book is *The Complete Guide to Windows Server 2008* (Addison-Wesley).

PROUD TO ANNOUNCE:
Recipient of the Eloqua "Marketing Center of Excellence" Award

Penton Marketing Services

WE KNOW YOUR CUSTOMERS

Now offering a full range of services that leverage our deep industry knowledge and customer relationships. From product launch to the final sale—put our years of experience to work for you.

Windows IT Pro

SQL Server magazine

SharePoint Pro CONNECTIONS

DevProConnections

MARKETING SERVICES WE OFFER:

- **Content Marketing:** Custom content that conveys technical information and thought leadership
- **SEO:** Optimizing your website to achieve higher search engine rankings
- **Video:** Leveraging the power of sight, sound, and motion to educate and inform
- **Social Media:** Monitoring, set-up and maintenance of social media vehicles
- **Mobile Marketing:** SMS messaging, WAP/App development and mobile marketing
- **Lead Lifecycling:** Lead generation, lead nurturing and tele-qualifying

FOR MORE INFORMATION:

PentonMarketingServices.com/tech
800 553 1945

Operations Manager Key Performance Indicators

Use the server health model to monitor your servers' performance

by Cameron Fuller

Monitoring your network involves a lot more than just keeping tabs on the health of your servers. It's also important to determine whether your websites, applications, network infrastructure, and servers are functioning 24 × 7. Using a single dashboard for network monitoring makes the task easier.

Before you start monitoring, it's useful to know what you're looking for in a healthy server. Server performance is typically assessed using four Key Performance Indicators (KPIs): Processor, Memory, Disk, and Network. We can create a health model for a server that incorporates these components, as well as other factors. For example, is a server healthy when it isn't running? If it's configured incorrectly? If its security is compromised? The more conditions we add to our definition of a healthy server, the more useful our health model will be in assessing our servers' health. A server's health model is sort of like a painting of what a server should look like—we start with a rough sketch of a server, adding details that help the sketch evolve into a full-color painting of the server.

Using the health model approach lets us provide monitoring not only for servers but also for custom applications, websites, network devices, and many other important aspects of a business. In Microsoft System Center Operations Manager 2007 R2, the server health model focuses on four main areas: availability, configuration, performance, and security. Several KPIs directly determine how well a server is performing—including Processor, Memory, Disk, and Network. The Windows Server Operating System Management Pack for Operations Manager 2007 includes the server health model. One of the methods for displaying Operations Manager's health model is the Health Explorer interface, which Figure 1 shows. This health model is extremely detailed; for the purposes of this article, let's focus on how Operations Manager integrates the various KPIs.

Processor KPI

Typically, a processor bottleneck is defined as more than 80 percent server utilization for a period of time. Unfortunately this type of bottleneck occurs relatively frequently and can generate a significant number of alerts that might not be actionable by the server. The Operations Manager monitor (Total CPU Utilization Percentage) takes processor monitoring a step further by alerting only when multiple conditions occur. Health states for this monitor are either healthy or critical based on the following conditions:

- Critical state occurs when processor utilization (Processor\% Processor Time\ _Total) is higher than 95 percent for 6 minutes (after three samples on a 2-minute schedule) *and* when processor queue length (System\Processor Queue Length) is greater than 15 for 4 minutes (after two samples on a 2-minute schedule).
- For all monitors discussed in this article, when the threshold decreases below the levels defined for the critical or warning state, the monitor resets itself to a healthy condition.

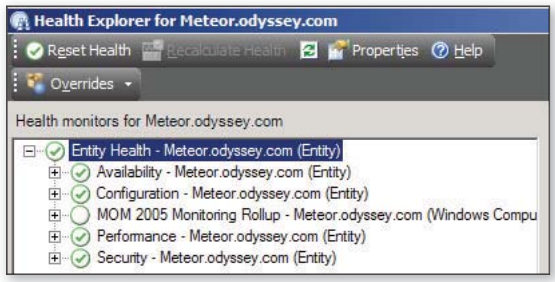

Figure 1: Health Explorer interface

This approach minimizes the amount of noise (i.e., nonactionable alerts) by providing an alert when the condition is likely to actually represent a bottleneck on a server versus a temporary spike in processor utilization. Although this approach provides a good starting point for most servers, not all servers are created alike. Some servers consistently experience higher processor workloads (e.g., servers running SQL Server or Exchange Server).

Virtualized servers also often experience higher than average processor interrupt levels that require tuning within Operations Manager. This doesn't indicate that the virtualized guest OS has additional overhead but rather that this particular counter might not be as relevant (or might have a higher than average value) in a virtualized guest OS.

Operations Manager lets you use overrides to tune alerts to detect different thresholds for different systems or groups of systems. An override changes the default behavior of a rule or monitor for the systems on which the override is applied. For example, suppose you have a computer group that contains all virtual servers. (For information about detecting both VMware and Hyper-V servers, see "Virtual Machine Discovery MP for Operations Manager 2007" at www.systemcentercentral.com/PackCatalog/PackCatalogDetails/tabid/145/IndexID/12572/Default.aspx.) You can target an override to change the thresholds to either a higher or lower level for that group. For processor counters, it's common to create an override that lowers the thresholds for systems on which processor bottlenecks are likely to occur or to increase the thresholds for the average processor interrupt level.

Operations Manager doesn't limit the health model for the processor to the Total CPU Utilization monitor. Instead, this monitor is supplemented with the Total Percentage Interrupt Time and the Total DPC Time Percentage counters. These counters can also indicate performance bottlenecks on the processor for a server.

The Total Percentage Interrupt Time monitor indicates the total percentage of interrupt time—which seems obvious; note that logical names are used throughout the health model, to let you easily determine which conditions a monitor evaluates. This monitor's healthy and critical states are defined as follows:

- Critical state occurs when the Total Percentage Interrupt Time monitor shows greater than 10 percent for 10 minutes (after five samples on a 2-minute schedule).

The Total DPC Time Percentage counter determines how much time the processor spends receiving and servicing deferred procedure calls, which are interrupts that run at a lower priority than standard interrupts. This monitor's healthy and critical states are defined as follows:

- Critical state occurs when the Total DPC Time Percentage monitor shows greater than 95 percent for 10 minutes (after five samples on a 2-minute schedule).

Operations Manager retains performance information in the operations database (called OperationsManager). This performance information can be used to create graphs for one or more systems. For example, Figure 2 shows processor utilization for several servers. This view is available in the Operations console. The Operations console reads directly from the OperationsManager database and can show data for up to the default retention period, which is 7 days.

Operations Manager also retains performance information in the data warehouse that can be used to provide trending of performance counters over time. Hourly and daily aggregated information is stored in the data warehouse for 400 days by default. Using the data warehouse lets us create reports that show performance information over a longer period of time than is available in the OperationsManager database.

Using the processor performance view or the processor performance report lets us establish a baseline for what the processor utilization looks like for a server or group of servers. We can then use this baseline to override the processor alerts to notify us when we've passed beyond what's considered normal behavior for a server's processor.

Operations Manager can perform diagnostic tasks, which provide information about what occurred when an object's health state changes, or recovery tasks, which repair the health state. A built-in diagnostic task occurs when a CPU goes from healthy to critical; the List Top CPU Consuming Processes diagnostic gathers information when a processor changes from a healthy to a critical state. This

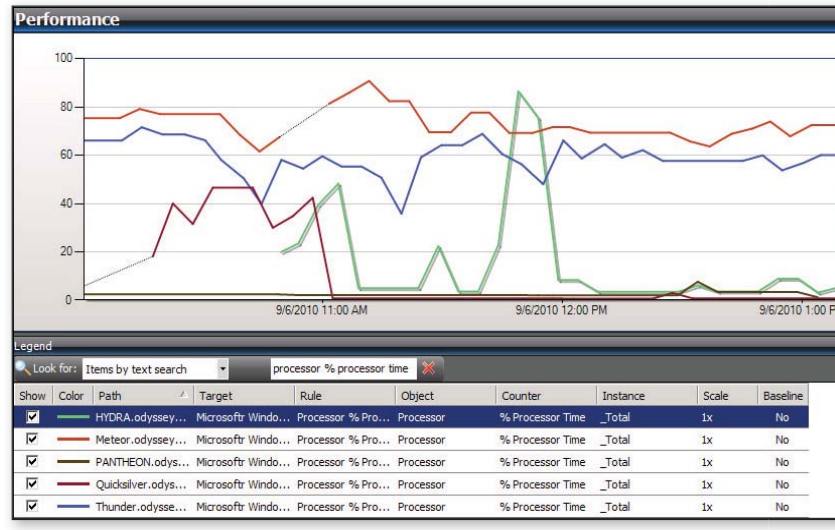

Figure 2: Processor utilization for several servers

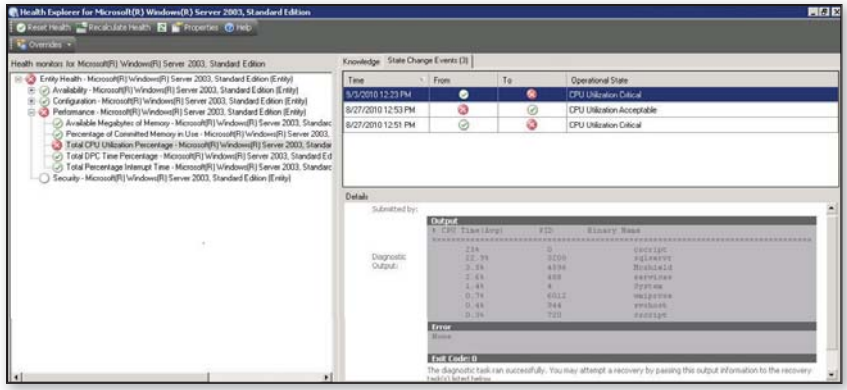

Figure 3: Viewing the built-in List Top CPU Consuming Processes diagnostic

information can be useful in determining why a server is experiencing a performance bottleneck. Figure 3 shows an example of this automated diagnostic.

Operations Manager provides a solid health model for what a processor's health should look like. This model is customizable on a per-object basis and is designed to be actionable. By combining both the processor health model and the ability to report trends in performance counters over a period of time, Operations Manager covers the Processor KPI extremely well.

Memory KPI

Memory bottlenecks are generally thought to exist when more than 80 percent of memory is committed on the server. The easiest solution is to simply add memory—but this solution is often not viable and sometimes not really necessary. Operations Manager tracks the percentage of committed and available memory and provides an alert when committed memory exceeds 80 percent (by default).

The Percentage of Committed Memory in Use monitor changes the server's health state based on the percentage of memory committed on the system. This monitor's healthy and critical states are defined as follows:

- Critical state occurs when committed memory is greater than 80 percent for 6 minutes (after three samples on a 2-minute schedule).

Operations Manager also monitors the amount of memory still available on a server. The Available Megabytes of Memory monitor changes the server's health state based on the number of available megabytes of memory on the system. This monitor's healthy and critical states are defined as follows:

- Critical state occurs when the available megabytes of memory falls below 2.5MB for 6 minutes (after three samples on a 2-minute schedule). By default, this value occurs only if a system is truly critical on memory. You might need to override the default value and set it to a larger number depending on your environment's requirements. Figure 4 shows performance monitoring for a server that's almost critical on memory but isn't yet close to the 2.5MB default threshold. To better use this monitor, you should create an override to increase the threshold from 2.5MB to a larger value based on the amount of memory on the server. According to the TechNet article "System Level Bottlenecks," at technet.microsoft.com/en-us/library/cc558658(BTS.10).aspx, a consistent value of less than 20 to 25 percent of installed RAM indicates insufficient memory.

The amount of paging is another aspect of memory monitoring that Operations Manager tracks as part of the Memory KPI. In the Windows Server 2008 Operating System Management Pack, this rule is called Memory Pages per Second 2008. Because it's a rule rather than a monitor, this rule doesn't affect the health model. However, Operations Manager does gather paging information to provide trending and potential bottlenecks for the OS.

The healthy state for these values varies depending on the types of applications that are installed on the server. For example, applications such as SQL Server and Exchange expand to use nearly all available memory on a system. In most environments, the administrator creates an override to set the memory threshold to between 95 and 99 percent for SQL Server systems and Exchange servers. To accomplish this task, you can use the SQL Server Management Pack's SQL Computers or SQL 2008 Computers groups and the Exchange Server Management Pack's Exchange 2007 Computer Group setting. For SQL Servers systems, you can implement a policy to restrict how much memory SQL Server can use; thus, thresholds can be based on an organization's SQL Server memory policy. In general, a threshold of 99 percent indicates a problem because it implies that the OS is most likely being starved for memory because of application memory requirements.

Just like the processor performance counters, the percentage of committed memory counter is available both in the operations database and in the data warehouse and can be used to provide trending information and to identify a baseline for normal memory utilization on a server. Figure 5 shows the percentage of committed memory for a server over a period of time. Like the processor health model, the

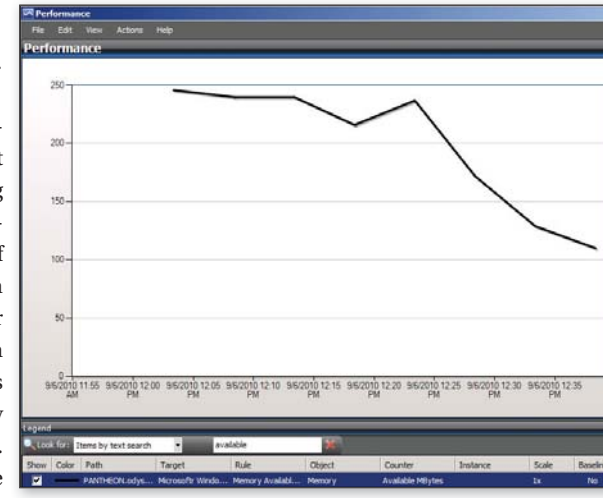

Figure 4: Server that's critical on memory but not at the 2.5MB threshold

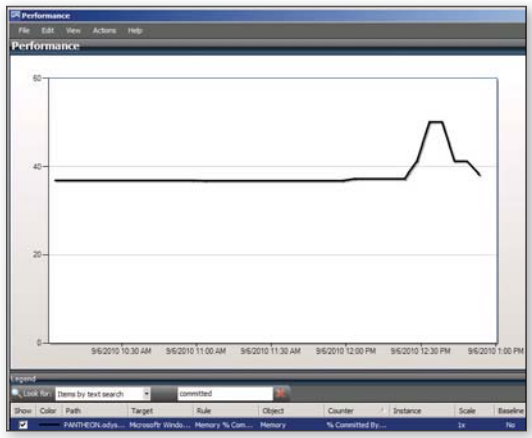

Figure 5: Percentage of committed memory for a server over a period of time

memory monitoring functionality within Operations Manager can be easily customized based on the requirements of an object or group of objects and provides another important piece of the overall health model for a server.

Disk KPI

Rather than being associated with too much reading and writing from the disk, disk bottlenecks are most often associated with how much free disk space exists on the drive. (Of course Operations Manager measures performance for both these metrics.)

To measure free disk space, Operations Manager tracks free megabytes and percentage of free space on all the drives on the servers it monitors. Operations Manager uses the Logical Disk Free Space monitor to determine the health of the disks it monitors. Like the processor monitor, two values determine when a drive is critical on free space: percentage of free space and megabytes of free space. The default values vary for drives depending on their function for the system. System drives (volumes that contain hardware-specific files needed to start Windows) have different thresholds than nonsystem drives because in general nonsystem drives are larger in size than system drives.

System drives are defined as healthy or in a warning or error state based on the following conditions:

- Warning state occurs when the percentage of free space is less than 10 percent *and* the actual free space is less than 200MB.
- Error state occurs when the percentage of free space is less than 5 percent *and* the actual free space is less than 100MB.

Nonsystem drives are defined as healthy or in a warning or error state based on the following conditions:

- Warning state occurs when the percentage of free space is less than 10 percent *and* the actual free space is less than 2,000MB.
- Error state occurs when the percentage of free space is less than 5 percent *and* the actual free space is less than 1,000MB.

Operations Manager also tracks how much the OS is reading and writing to the disk, as well as current disk queue length. In addition, health monitoring is available for disk transfers (Average Disk Seconds Per Transfer) and fragmentation level of a drive. Monitoring of disk reads (Average Disk Seconds Per Read) and disk writes (Average Disk Seconds Per Write) is disabled by default but can be enabled through an override.

Disk utilization is determined by the Average Disk Seconds Per Transfer monitor. This monitor's healthy and critical states are defined as follows:

- Critical state occurs when the average disk seconds per transfer is greater than 50 for 5 minutes (after five samples on a 1-minute schedule).

Fragmentation health is determined by the Logical Disk Fragmentation Level monitor. This monitor's healthy and warning states are defined as follows:

- Warning state occurs when the percentage of file fragmentation is greater than 10 percent. (This monitor checks health state once a day at 3:00 A.M. on Saturday by default.)

The Logical Disk Fragmentation Level monitor also includes a recovery task called Logical Disk Defragmentation, which is disabled by default. This task can automatically run a defragmentation if the drive exceeds the threshold defined for the monitor. (For more information about this monitor, see "OpsMgr ReSearch This KB – Logical Disk Fragmentation Level is High" at blogs.catapultsystems.com/cfuller/archive/2010/05/19/opsmgr-research-this-kb-%E2%80%93-logical-disk-fragmentation-level-is-high.aspx.)

Operations Manager also checks the availability of logical disks on a system every 5 minutes through the Logical Disk Availability monitor. This monitor provides an alert if a drive disappears or becomes inaccessible by a server that Operations Manager monitors. In most cases this monitor functions well as designed. However, I ran into a situation in which a volume was mounted and dismounted on a scheduled basis for a server. In this case I had to create an override to disable the monitor for that drive.

The Disk KPI is fully covered by the Operations Manager health model. Free space and availability are monitored, and health state is defined—including the amount of data being transferred and even the fragmentation level of the drive.

Network KPI

Operations Manager tracks performance information for network adapters through the following three counters:

- Bytes Received/sec
- Bytes Sent/sec
- Bytes Total/sec

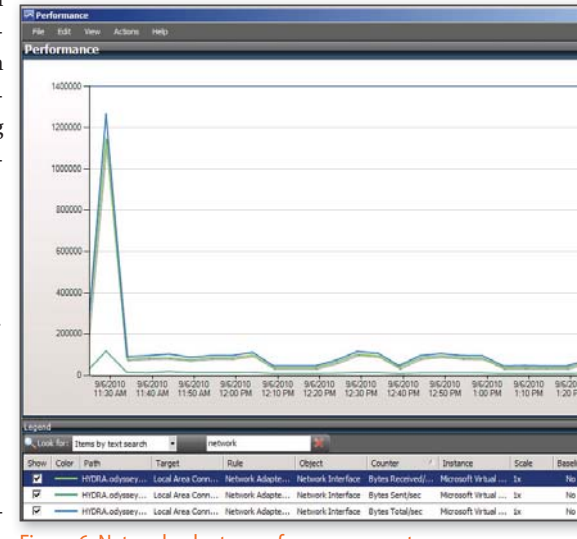

Figure 6: Network adapter performance counters

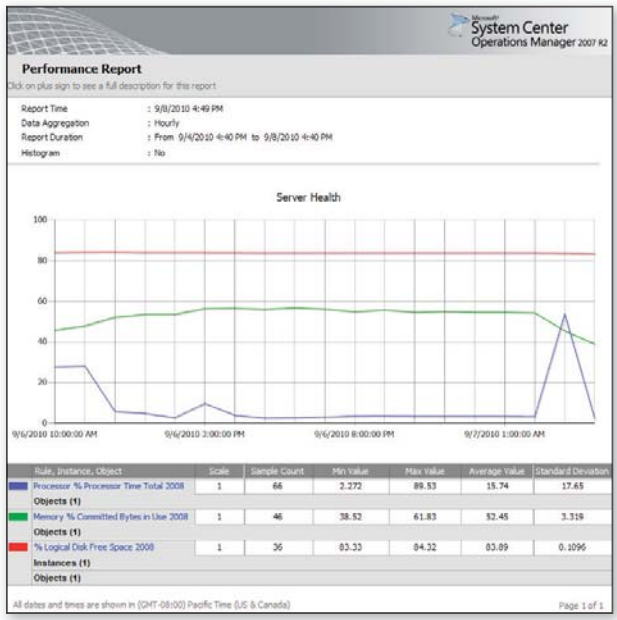

Figure 7: Custom report showing three different counters

These counters are tracked by default. You can display them through the Operations Manager console or reports. Figure 6 shows the network adapter performance counters. These performance counters don't have associated monitors, so they don't affect the network adapter health model.

Operations Manager provides a monitor called Network Adapter Connection Health that can change the health state of a network adapter if it's removed from a server. This monitor is disabled by default (probably to minimize nonactionable alerts generated by servers in an environment with multiple disconnected network adapters); it can be enabled to provide availability information as part of the Operations Manager health model for servers.

Reporting

Operations Manager includes a reporting function based on SQL Server Reporting Services (SSRS). Various Operations Manager management packs provide built-in reports that use the SSRS functionality. Figure 7 shows a custom report that was created in only a few minutes using processes discussed in "Creating Useful Custom Reports in OpsMgr: Gathering Custom Performance Counters" (blogs.catapultsystems.com/cfuller/archive/2010/07/21/creating-useful-custom-reports-in-opsmgr-gathering-custom-performance-counters.aspx). This report shows a single server with three different counters (percent processor time, percent committed bytes, and percent logical disk free space on the C drive).

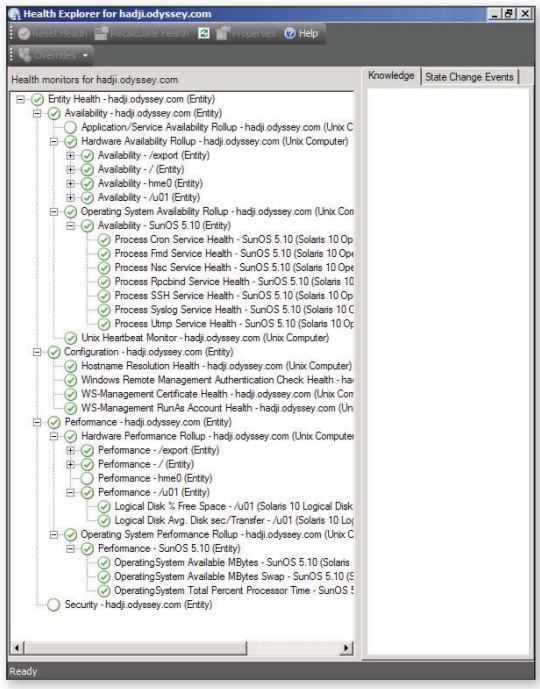

Figure 8: Operations Manager integration with UNIX

Non-Windows Systems

With the release of Operations Manager 2007 R2, Operations Manager includes the ability to monitor UNIX and Linux systems with an open-source agent that's deployed to the system(s). Figure 8 shows how Operations Manager integrates the components from a UNIX system into the health model in a similar method to how a Windows server is integrated. Note the Processor, Memory, and Disk KPIs.

For more information about Operations Manager 2007 R2's cross-platform integration, see *System Center Operations Manager (OpsMgr) 2007 R2 Unleashed* (Sams, 2010). For information about cross-platform processor monitoring support, see "Understanding CPU Performance Counters on Cross Platform Monitors" at blogs.msdn.com/b/scxplat/archive/2010/02/04/understanding-cpu-performance-counters-on-cross-platform-monitors.aspx.

The Big Picture

Operations Manager's health model lets us take basic concepts that have historically determined servers' health, such as "Is the processor running at more than 80 percent?" and expand them into a more comprehensive and customizable model to better evaluate servers' health. The health model maps out how well servers are performing from an OS level, which is a cornerstone for server health. A server's OS health model, combined with an application's health model and other models, creates a much better picture of the server and therefore the environment as a whole.

Operations Manager uses health models that range from the lowest level of an object, such as a disk or processor, to far more complicated structures, such as a distributed application like Exchange or Active Directory (AD). A good analogy is to think of health models as building blocks; Operations Manager uses these blocks to build larger structures, such as a custom-built geographically dispersed application.

Distributed applications in Operations Manager use health models to perform the same tasks as for a server but on a larger scale. These distributed applications can then be incorporated into a dashboard solution that lets you simultaneously monitor your websites, applications, network infrastructure, and servers.

InstantDoc ID 128969

Cameron Fuller
(cameron.fuller@catapultsystems.com) is principal consultant for Catapult Systems, an IT consulting company and Microsoft Gold Certified Partner. He's an Operations Manager MVP and co-author of *System Center Operations Manager 2007 Unleashed* and *System Center Operations Manager 2007 R2 Unleashed* (Sams).

FEATURE

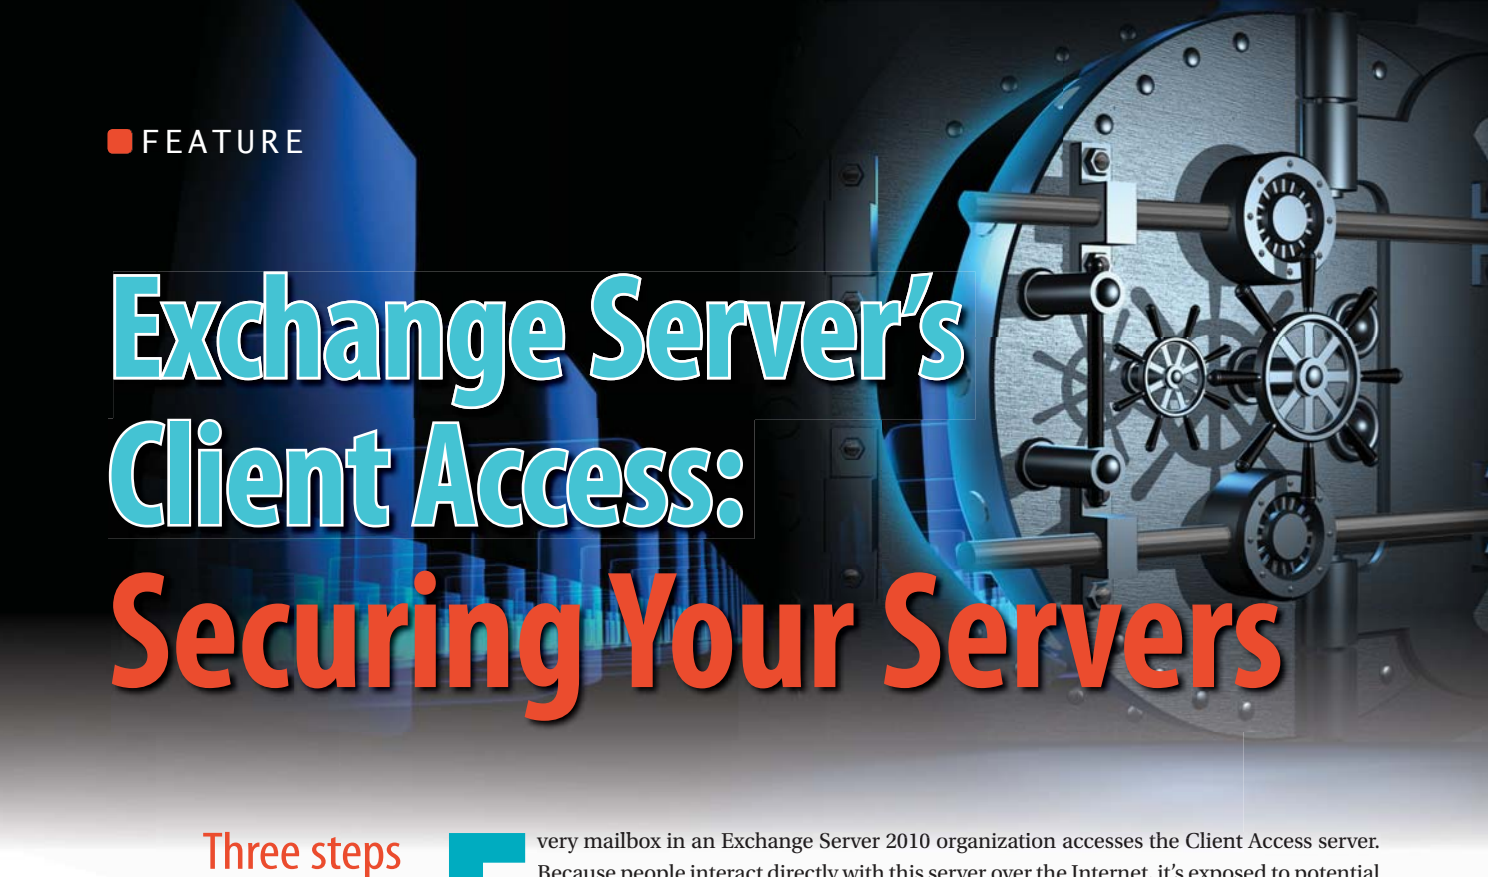

Exchange Server's Client Access: Securing Your Servers

Three steps for greatly increased security

by Ken St. Cyr

Every mailbox in an Exchange Server 2010 organization accesses the Client Access server. Because people interact directly with this server over the Internet, it's exposed to potential attackers. Therefore, ensuring the Client Access server's security is paramount. In previous articles I walked you through various aspects of the Client Access server role and showed you how to deploy it and load-balance it. (See the Learning Path for previous articles about the Client Access server.) In this article I discuss how to secure the Client Access server.

The term security means different things to different people. We often think of security as an ambiguous process that's used to determine how accessible a server is. But when considering security for Client Access servers, there are three steps you can take to greatly increase security:

1. Use certificates for external services.
2. Harden the server OS.
3. Limit exposure with a reverse proxy.

Use Certificates for External Services

Clients access data through Client Access servers in various forms. For email, this access might include connecting via Outlook Web App (OWA), ActiveSync, Outlook Anywhere, Exchange Web Services (EWS), POP, IMAP, or MAPI on an Outlook client. All these methodologies and protocols, except for MAPI on an Outlook client, were designed to be used over the public Internet. You must therefore ensure that the data being transmitted over the Internet can't be read by people who might be capturing packets on the network. In Exchange, you can accomplish this by encrypting these publicly routable connections with SSL certificates.

An SSL certificate has a pair of keys associated with it; one key is private and the other is public. These keys are used for encrypting the data passed between two parties. The server that receives the certificate (the Client Access server in this case) holds a copy of the private key. The public key is given to the clients that access the Client Access server. The data that's encrypted with the public key can be decrypted only with the private key. This allows the client to encrypt the data it sends to the Client Access server and ensures that only the Client Access server can decrypt it. However, you also need to ensure that the data the Client Access server sends back to the client is encrypted as well. To accomplish this task, the client generates a shared key that can be used to both encrypt and decrypt the data. The client then encrypts that shared key with the public key of the Client Access server and sends it to the Client Access server. The Client Access server and the

> You need to ensure that the data the Client Access server sends back to the client is encrypted.

Learning Path

***WINDOWS IT PRO* RESOURCES:**

"Exchange Server's Client Access: An Introduction," InstantDoc ID 125061

"Exchange Server's Client Access: Deploying Your Servers," InstantDoc ID 125347

"Exchange Server's Client Access: Load Balancing Your Servers," InstantDoc ID 125863

client then have a shared key that they can both use to encrypt and decrypt the data. This prevents other people, who might be capturing packets on the network, from reading the data transmitted between the client and the Client Access server. Even if someone captures the data, they need the shared key to decrypt it. And because the shared key is held only by the client and the Client Access server, they're out of luck.

The first thing you need to do to make your data transmission secure is obtain a certificate for your Client Access servers. When you install the Client Access server, a self-signed certificate is created by default. This certificate is generated by the Client Access server itself. The problem with using a self-signed certificate is that your clients don't trust it. If your Client Access server is using a self-signed certificate and you try to access OWA, you'll receive an error message similar to the one in Figure 1. This doesn't necessarily mean that your Client Access server isn't secure; it simply means that your clients don't trust the certificate issuer (i.e., the Client Access server). If you add the self-signed certificate to the list of trusted issuers on the client, the problem is solved. However, this is difficult to manage in a distributed environment. This is particularly true in Exchange Server 2007 because the default self-signed certificate is good only for a year. After the certificate expires, you have to renew it and redistribute it. Exchange 2010 extends the expiration period to five years. Even so, self-signed certificates aren't an ideal solution.

A better solution is to obtain a mutually trusted certificate from a valid Certificate Authority (CA). The idea is that the Client Access server and the client both trust the certificate, so they both believe that the certificate is valid. You can obtain such a certificate from your own CA (if you have one) or you can buy a certificate from a third party, such as DigiCert, VeriSign, Entrust, or GoDaddy.

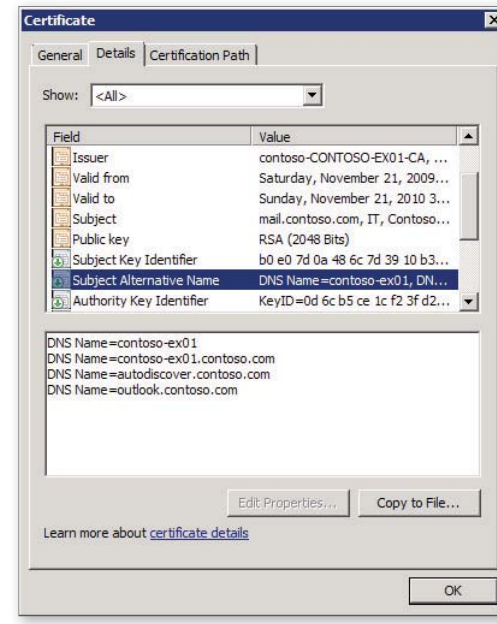

Figure 2: SAN certificate alternative names

The first thing you need to do to make your data transmission secure is obtain a certificate for your Client Access servers.

Client Access servers can use many different names to communicate. Therefore you need to ensure that all the names used to access the server are listed in the certificate when you obtain it. If not, your server will fail the identity check. Two types of certificates let you specify more than one name for a server: a Subject Alternative Name (SAN) certificate and a wildcard certificate.

Wildcard certificates can be expensive, but they let you use a wildcard character for the certificate's subject name. Thus, any name by which the server is accessed is considered valid. SAN certificates are cheaper and more commonly used. A SAN certificate lets you list the server's primary name in the Subject Name field and alternative names in the Subject Alternative Name field, as Figure 2 shows.

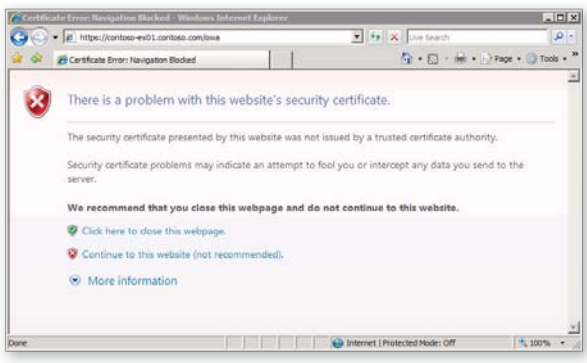

Figure 1: Security certificate error message

To obtain a certificate, you must generate a certificate request file on the Exchange server and give that file to the CA that will issue the certificate. When the request file is created, the private key portion of the certificate is created and stored on the server that generated the request. This keeps the private key secure because you don't have to give it to the CA—so your server is the only party that has it. You can generate this certificate request in Exchange 2010 in one of two ways: by using Exchange Management Console's (EMC's) New Exchange Certificate wizard or by using Exchange Management Shell's (EMS's) New-ExchangeCertificate cmdlet.

To launch the New Exchange Certificate wizard in EMC, select the Client Access server you're requesting the certificate for and select New Exchange Certificate, as Figure 3 shows. In the New Exchange Certificate wizard, you have the option of requesting certificates for some or all of the services running on the Client Access server, as Figure 4 shows.

After you complete this process, you'll have a certificate request file that you can give to your CA. The CA uses this file to generate a certificate. After you receive the certificate from the CA, you must then import the certificate into Exchange to complete the process. To do so, start EMC and select the certificate request that you generated for the Client Access server. Select Complete Pending Request,

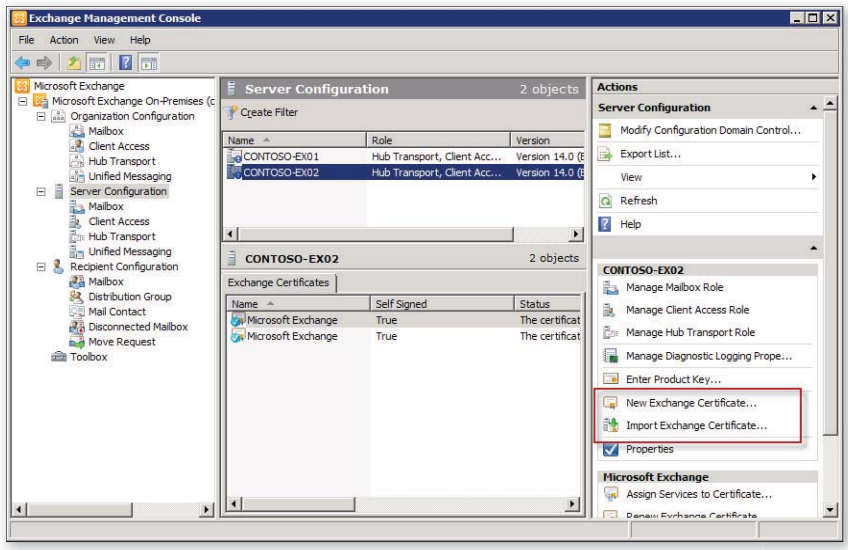

Figure 3: Launching the New Exchange Certificate wizard in EMC

as Figure 5 shows, to give Exchange the certificate and enable it.

A question that's often asked is whether every Client Access server can use the same certificate. For example, in a Client Access array, every server is accessed using the same external name—so, can you import the same certificate on every server, or do you need a separate certificate for each server? The answer depends on a couple of factors. If you purchase your certificate from a publically trusted CA, such as Thawte, then you must pay for each certificate that you want to use on each Client Access server. But if you use an internal CA, you aren't legally obligated to use multiple certificates. In this case you can use a single wildcard certificate. A typical Exchange certificate often uses multiple SANs, so if you use the same certificate for multiple Client Access servers, you should make sure that the valid names for each server are contained in the Subject Alternative Name field.

Harden the Server OS

Hardening a server involves removing extraneous services and processes that could be potential sources of vulnerability. Even if no known vulnerability exists, simply disabling components that you don't use helps ensure protection if a security hole is found in the future. Exchange 2007 includes templates that you can use with the Windows Security Configuration Wizard (SCW) to lock down Exchange on a role-by-role basis. Exchange 2010 doesn't include these templates because each role is locked down by default. However, this doesn't mean the OS is locked down. Some basic steps you can take to harden Windows Server 2008 R2 or Server 2008 include the following:

- Design your organizational unit (OU) structure for role-based policies.
- Implement the security policies identified in the Enterprise Client (EC) settings or the Specialized Security – Limited Functionality (SSLF) settings.
- Audit the security logs on your server.

The primary method for hardening Client Access servers is to use Group Policy Objects (GPOs) to enforce policies on the servers. To make this process work efficiently, you need to ensure that your OU structure is configured to let you lock down your Exchange servers with a common policy and then lock down individual servers by role. When you use this method of GPO enforcement, you independently secure your Client Access role from other server roles. For example, you can disable the POP or IMAP service for all Exchange servers but enable it for Client Access servers. Figure 6 illustrates how this methodology of GPO application works. For additional information on this topic, see the TechNet article "Designing OU Structures that Work" (technet.microsoft.com/en-us/magazine/2008.05.oudesign.aspx).

After your OU structure is in place, you can implement Group Policy settings to decrease the attack surface of your servers. You'll want to start by implementing a baseline security configuration that you can later build on. Two baselines are worth mentioning, but only one of them is reasonable to use in most situations.

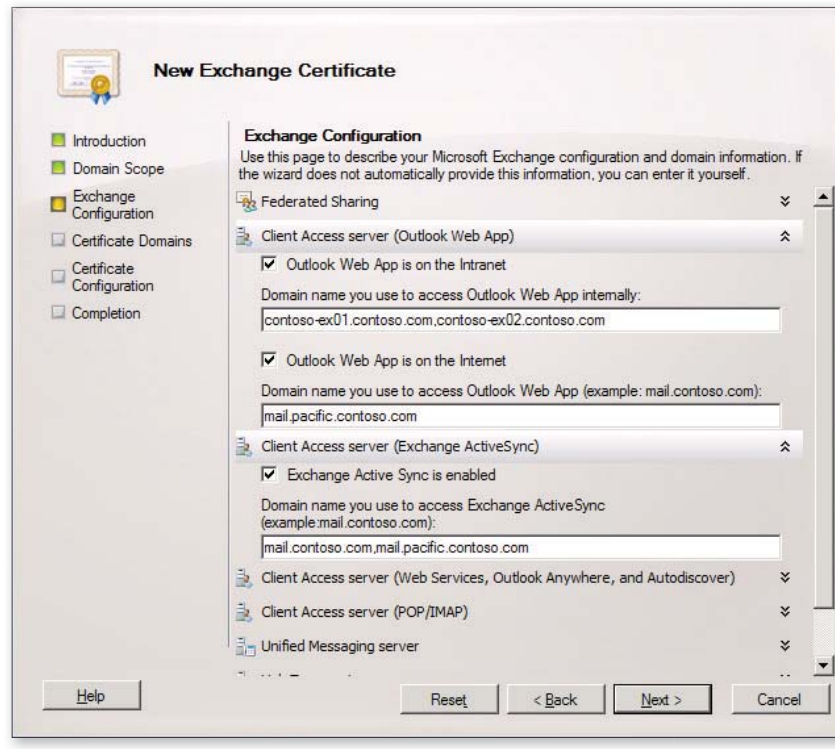

Figure 4: Requesting certificates for specific services running on the Client Access server

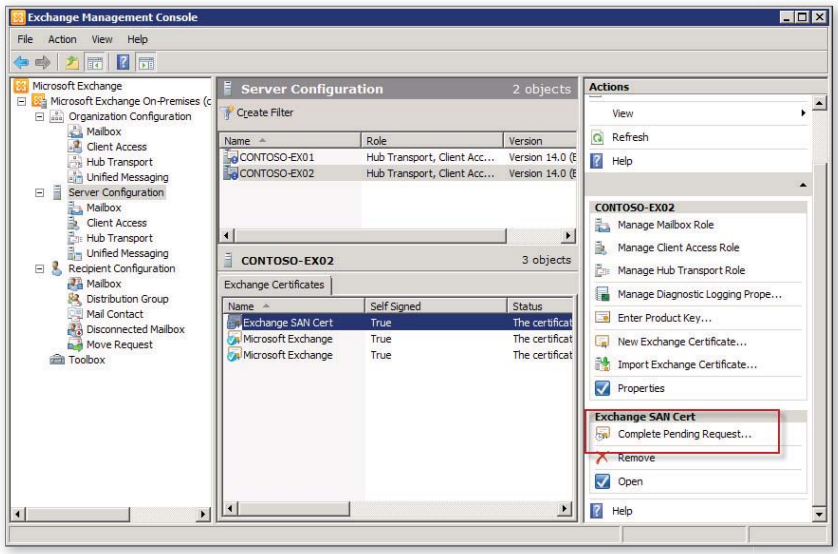

Figure 5: Completing a certificate request

The EC baseline includes some solid and well-accepted security settings. (Don't let the name fool you; this baseline applies to Windows Server 2003 and later member servers as well as client OSs.) EC baseline security settings include disabling the guest account and disabling cached credentials.

The SSLF baseline is tightly locked down. (Note the "limited functionality" portion of its name.) SSLF goes to extremes in securing a server—so much so that you'll likely need to troubleshoot some problems and roll back to pre-baseline policy settings. I've seen some odd behavior with SSLF, such as EMC connectivity experiencing intermittent failures.

I recommend using the EC baseline rather than SSLF for securing your Exchange servers. You can download this baseline, as well as documentation for implementing the settings, as part of the Microsoft Security Compliance Manager at technet.microsoft.com/en-us/library/cc677002.aspx.

A final note about hardening your OS is to pay attention to the audit logs that are generated. If you use the security baselines that I recommend, then security auditing is configured on the server. But configuring auditing is just the first step. You also need a method of regularly monitoring the audit logs in order to detect intrusion attempts. Many tools are available to consolidate your servers' security audit logs. The two most common tools I've seen used are Microsoft System Center Operations Manager's Audit Collection Services (ACS) and NetIQ Security Manager. Each product has unique features, but the idea is the same—collect your security logs and pay attention to them.

Limit Exposure with a Reverse Proxy

Perhaps the most important aspect of protecting your Client Access servers is to ensure that people who aren't part of your organization can't break into them. Limiting exposure to your Client Access servers means keeping them within your protected environment but still making them accessible by legitimate users over the Internet. Not only is this the preferred configuration, but keeping your Client Access servers out of your demilitarized zone (DMZ) is also the only configuration that Microsoft supports. Instead of publishing your Client Access server directly to the Internet, you publish a reverse proxy server to the Internet and the reverse proxy subsequently provides connectivity to the Client Access server. Figure 7 illustrates how this scenario takes place.

Numerous reverse proxy server products are available, including Microsoft's own Internet Security and Acceleration (ISA) Server, Forefront Threat Management Gateway (TMG—ISA Server's 64-bit successor), Apache, Squid, various flavors of Linux with strict lockdowns, and enterprise-class hardware devices such as F5's BIG-IP. Each reverse proxy server offers different advantages and disadvantages, but the basic concept is the same. The reverse proxy server's job is to intelligently accept or reject connections coming from the Internet and pass them back to the Client Access server. The level of intelligence that the reverse proxy server has varies between products. Some reverse proxy servers provide access based on simple rules, whereas others have a deep knowledge of the applications they're protecting. The advantage of these more intelligent reverse proxy servers is that they know what kind of traffic the application will or won't accept, so they can block application requests that are malformatted or obvious hack attempts. You need to keep a few considerations in mind when configuring a reverse proxy server to protect your Client Access servers, such as how authentication is performed, how SSL connectivity is handled, and whether you'll use the reverse proxy as a load balancer.

One of the advantages of a reverse proxy is the ability to perform preauthentication. When you configure preauthentication, you let the reverse proxy authenticate the user before the request is passed back to the Client Access server. Authentication still occurs on the Client Access server, but you can delegate that authentication so that the user doesn't need to re-enter any credentials. For example, you can implement forms-based authentication on your reverse proxy and implement NTLM authentication on your OWA virtual directory on the Client Access server. In this case, when an Internet user accesses OWA,

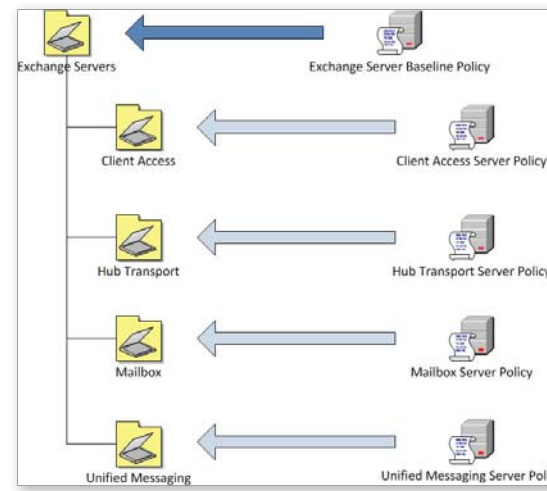

Figure 6: Using GPOs to harden Client Access servers

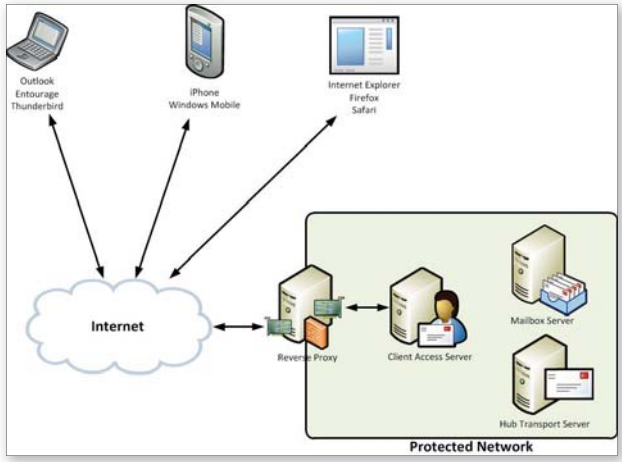

Figure 7: Using a reverse proxy server to limit your Client Access servers' exposure

he or she is prompted for forms-based authentication through the preauthentication process. However, when a user that's on your internal network accesses OWA directly on the Client Access server, his or her NTLM credentials are passed right through and the user isn't prompted for authentication. The thing you need to be careful about with preauthentication is that it isn't supported by all email access methods. For example, if you require preauthentication for Outlook Anywhere clients, they won't connect because Outlook doesn't know how to handle the preauthentication request. Also, if you're using Exchange 2010's federated sharing features, then you need to exclude the following paths from preauthentication:

- /EWS/exchange.asmx/wssecurity
- /autodiscover/autodiscover.svc
- /autodiscover/autodiscover .svc/wssecurity

The primary thing to keep in mind for preauthentication is that it's good to use in scenarios in which email access will be browser based, such as in the case of OWA and the Exchange Control Panel (ECP).

When it comes to SSL connectivity, you typically have a few choices available, depending on the reverse proxy that you use. In general these options are:

- Basic firewall: Maintain the same SSL connection from the client to the Client Access server.
- SSL offloading: Terminate the SSL connection at the reverse proxy and use an unencrypted connection from the reverse proxy to the Client Access server.
- SSL bridging: Terminate the SSL connection at the reverse proxy and establish a new SSL connection from the reverse proxy to the Client Access server.

These options each have benefits. In the first option, the reverse proxy acts like a basic firewall and just passes the connection straight through to the Client Access server, in a fashion similar to port forwarding. This method tends to be the least secure because the reverse proxy doesn't perform any content inspection. The second option, SSL offloading, lets the client establish an SSL connection with the reverse proxy, but the connection from the reverse proxy to the Client Access server is unencrypted. The advantage with this method is that because the reverse proxy performs the decryption and verification associated with the SSL connection, the Client Access server doesn't have to. The third option, SSL bridging, is the most secure. When SSL bridging is used, the connection is terminated at the reverse proxy so that the reverse proxy can perform inspection. Then the reverse proxy reestablishes a new SSL connection back to the Client Access server in order to keep the session encrypted. In this option, the SSL termination for the client occurs at the reverse proxy and the SSL termination for the reverse proxy occurs at the Client Access server, as Figure 8 shows.

A final consideration to keep in mind is whether to configure your reverse proxy as your load balancer for your Client Access servers. This was a better option in previous versions of Exchange than in Exchange 2010 because Exchange 2010's MAPI clients interact directly with Client Access servers. Therefore, you need to ensure that the MAPI clients connect to the load-balanced array of Client Access servers instead of an individual Client Access server. (For more information about load balancing, see "Exchange Server's Client Access: Load Balancing Your Servers," InstantDoc ID 125863.) You can make this scenario work, but not without some pain. You must route the connections for your internal users through the external interface and back into the network, which means you need to allow RPC connections through your reverse proxy—which isn't a good idea.

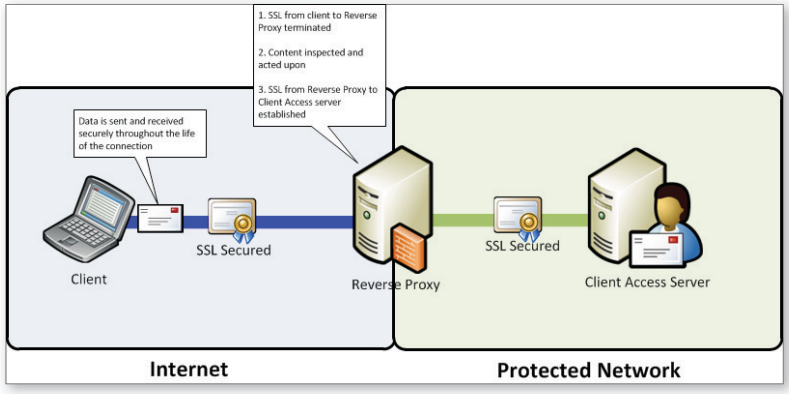

Figure 8: SSL bridging

Easier than Expected

Securing your Client Access servers isn't as difficult a task as you might imagine. Exchange 2010's built-in security options make the process simple; all you need to do is ensure that you have the right certificates, that your server OS is locked down, and that you're using a reverse proxy to protect your servers from the Internet.

InstantDoc ID 128939

Ken St. Cyr
(ken.stcyr@microsoft.com) is a solution architect at Microsoft with more than 10 years of industry experience. He's a Microsoft Certified Master in Directory Services and the author of *Exchange Server 2010 Administration Instant Reference* (Sybex).

“THE CONVERSATION BEGINS HERE”

Microsoft® Exchange CONNECTIONS

WINDOWS CONNECTIONS

UNIFIED COMMUNICATIONS CONNECTIONS

SharePoint CONNECTIONS

SQL Server CONNECTIONS

COLOCATED WITH THESE EXCITING EVENTS:

Microsoft® Visual Studio CONNECTIONS

Microsoft® ASP.net CONNECTIONS

Microsoft® Silverlight CONNECTIONS

BONUS:
Mobile Apps Track

QUESTIONS ANSWERED • STRATEGY DEFINED • RELATIONSHIPS BUILT

One Place, One Time...

Make **CONNECTIONS** the **CONFERENCE** you bring your whole team to this year!

MARCH 27-30, 2011

ORLANDO, FL

GRANDE LAKES JW MARRIOTT RESORT HOTEL

REGISTER EARLY!

Reserve a room by December 20th *for a minimum of 3 nights*
and receive a **$100** Marriott gift certificate,
plus a **$100** discount from the regular registration fee.

WinConnections ... Providing the vision + intelligence to keep you and your company competitive in today's market!

Only Microsoft and Industry Experts speak at WinConnections!

KEYNOTES AND INDUSTRY EXPERTS

QUENTIN CLARK
MICROSOFT

STEVE FOX
MICROSOFT

SCOTT GUTHRIE
MICROSOFT

MARK MINASI
MINASI RESEARCH AND DEVELOPMENT

TONY REDMOND
TONY REDMOND AND ASSOCIATES

PAUL THURROTT
WINDOWS IT PRO

CHECK WEB SITE FOR DESCRIPTIONS OF SESSIONS AND WORKSHOPS

www.WinConnections.com • 800-438-6720 • 203.400.6121 • Register Today!

Microsoft® SharePointPro CONNECTIONS SQL Server WindowsITPro TECH Conferences PENTON MEDIA

WINDOWS CONNECTIONS SESSIONS

WIN01: DON JONES' 75-MINUTE POWERSHELL CRASH COURSE
DON JONES

WIN02: DON JONES' SECRETS OF CLIENT AND SERVER REMOTE CONTROL WITH WINDOWS POWERSHELL
DON JONES

WIN03: DON JONES' ADVANCED WINDOWS POWERSHELL: ERROR HANDLING, DEBUGGING, "SCRIPT CMDLETS," AND MORE
DON JONES

WIN04: VDI-IN-A-BOX: MICROSOFT DESKTOP VIRTUALIZATION FOR SMALLER SCENARIOS AND BUSINESSES
GREG SHIELDS

WIN05: PREPARING SOFTWARE FOR DEPLOYMENT WITH A WINDOWS 7 UPGRADE
GREG SHIELDS

WIN06: AUTOMATICALLY DEPLOYING WINDOWS 7 WITHOUT THE MICROSOFT ALPHABET SOUP
GREG SHIELDS

WIN07: MICROSOFT OPALIS 101: YOUR SECRET IT PRO AUTOMATION BUDDY, ENGINE, AND SECRET WEAPON
JEREMY MOSKOWITZ

WIN08: MICROSOFT AND 3RD-PARTY GPO TOOLS YOU HAVE NEVER HEARD OF (AND SOME YOU HAVE)
JEREMY MOSKOWITZ

WIN09: TOTAL WORKSTATION LOCKDOWN: YOUR ACTION PLAN
JEREMY MOSKOWITZ

WIN10: VMWARE ESX BEST PRACTICES: NOTES FROM THE FIELD
ALAN SUGANO

WIN11: THE CLOUD CONTROVERSY: AN IN-THE-TRENCHES VIEW OF YOUR COMPANY'S PLACE IN THE CLOUD
ALAN SUGANO

WIN12: PASS A PAYMENT CARD INDUSTRY (PCI) COMPLIANCE SCAN (AND WHY YOU'D WANT TO EVEN IF YOU DON'T HAVE TO)
ALAN SUGANO

WIN13: CONDUCTING A FORENSIC COMPUTER INVESTIGATION FOR IT STAFF
MIKE DANSEGLIO

WIN14: NETWORK SNIFFING FOR IT PROS, NOT HACKERS
MIKE DANSEGLIO

WIN15: THE NETWORK FILES, CASE #53: DIAGNOSING DISEASES OF DNS
MARK MINASI

WIN16: BEND R2'S ACTIVE DIRECTORY TO YOUR WILL
MARK MINASI

WIN17: TEN (OR MORE) THINGS YOU PROBABLY DON'T KNOW ABOUT WINDOWS SERVER 2008 R2
MARK MINASI

WIN18: GOING, GOING, GONE? VIRTUALIZING YOUR ACTIVE DIRECTORY FOREST
SEAN DEUBY

WIN19: THE BEST FREE TOOLS FOR WINDOWS DESKTOP ADMINISTRATION
GREG SHIELDS

WIN20: ACTIVE DIRECTORY FEDERATION SERVICES (ADFS) – WHY YOU SHOULD CARE AND WHAT YOU SHOULD KNOW
LAURA HUNTER

WIN21: INSTALLING ACTIVE DIRECTORY FEDERATION SERVICES (ADFS) 2.0
SEAN DEUBY

WIN22: HOW WIN IT DEPLOYED ACTIVE DIRECTORY FOR WINDOWS SERVER 2008 AND R2
LAURA HUNTER

WIN23: BETTER WINDOWS IMAGING: THE VIRTUAL HARD DISK (VHD) FORMAT
RHONDA LAYFIELD

WIN24: DEPLOYING WINDOWS IMAGES THE SAFE, SECURE WAY
RHONDA LAYFIELD

EXCHANGE CONNECTIONS SESSIONS

EXC01: TELEPHONY DEMYSTIFIED FOR EXCHANGE ADMINS (PART 1)
PETER O'DOWD

EXC02: TELEPHONY DEMYSTIFIED FOR EXCHANGE ADMINS (PART 2)
PETER O'DOWD

EXC03: THE EXCHANGE SERVER STORE DEMYSTIFIED
PETER O'DOWD

EXC04: CAS 2010 - MORE FOOD FOR THOUGHT
KEVIN LAAHS

EXC05: FEAR WEB SERVICES NO MORE
KEVIN LAAHS

EXC06: EXCHANGE, SHAREPOINT AND OFFICE - BETTER TOGETHER?
KEVIN LAAHS

EXC07: THE RPC CLIENT ACCESS ARRAY: THE MISSING PIECE OF EXCHANGE AVAILABILITY
DEVIN L. GANAGER

EXC08: WAN OPTIMIZATION FOR EXCHANGE
DEVIN L. GANGER

EXC09: LOAD BALANCING YOUR EXCHANGE DEPLOYMENT
DEVIN L. GANGER

EXC10: EXCHANGE 2010 HIGH AVAILABILITY WITHOUT THE HIGH COST
JIM MCBEE

EXC11: MIGRATING TO EXCHANGE 2010 FROM EXCHANGE 2003
JIM MCBEE

EXC12: MAKING GOOD IT BUSINESS DECISIONS WHILE CLOUD PROOFING YOUR CAREER
JIM MCBEE

EXC13: FOREFRONT TMG CLIENT ACCESS PUBLICATION AND EDGE TRANSPORT INTEGRATION
MIKE CROWLEY

ABSTRACTS ARE AVAILABLE ONLINE

EXC14: INFORMATION RIGHTS MANAGEMENT EXPLORED
MIKE CROWLEY

EXC15: OFFICE 365
MIKE CROWLEY

EXC16: HIGH-AVAILABILITY WITH THE OTHER ROLES: HUB TRANSPORT, CLIENT ACCESS, AND UNIFIED MESSAGING
MICHAEL B. SMITH

EXC17: DUMPSTER AND LITIGATION HOLD – DUMPSTER 2.0 VS. DUMPSTER 1.0
MICHAEL B. SMITH

EXC18: CONFIGURATION AND USAGE OF RETENTION POLICIES IN EXCHANGE 2010 SP
MICHAEL B. SMITH

EXC19: EXCHANGE 2010 DEPLOYMENT AND MIGRATION BEST PRACTICES
KIERAN MCCORRY

EXC20: EXCHANGE 2010 INFORMATION PROTECTION AND RETENTION
KIERAN MCCORRY

EXC21: EXCHANGE 2010 SERVICE PACK 1
KIERAN MCCORRY

MICROSOFT SESSIONS

HOW MICROSOFT IT IMPLEMENTED MICROSOFT EXCHANGE SERVER 2010

WHAT'S NEW IN ARCHIVING, RETENTION, AND DISCOVERY IN MICROSOFT EXCHANGE SERVER 2010 SP1

WHAT'S NEW IN OWA, MOBILITY AND CALENDARING IN MICROSOFT EXCHANGE SERVER 2010 SP1

MICROSOFT® LYNC™ SERVER 2010: TRANSFORMING THE WAY PEOPLE COMMUNICATE

MICROSOFT® LYNC™ SERVER 2010: WHAT'S NEW IN DEVICES

BUILDING COMMUNICATIONS ENABLED BUSINESS PROCESSES WITH MICROSOFT® LYNC™ SERVER 2010

MICROSOFT® LYNC™ SERVER 2010 INTEROPERABILITY: VOICE, VIDEO, CONFERENCING, IM, AND PRESENCE

SHAREPOINT CONNECTIONS SESSIONS

DEVELOPMENT TRACK

HDEV01: DEVELOPERS DEEP DIVE INTO SHAREPOINT SECURITY
TED PATTISON

HDEV02: SHAREPOINT DATA ACCESS SHOOTOUT
TED PATTISON

HDEV03: ADVANCED CONTROL AND WEB PART DEVELOPMENT
TED PATTISON

HDEV04: RECORDS MANAGEMENT IMPROVEMENTS IN SHAREPOINT 2010
JOHN HOLLIDAY

HDEV05: SHAREPOINT 2010 RECORDS MANAGEMENT DEVELOPMENT
JOHN HOLLIDAY

HDEV06: CONTENT TYPE DISCOVERY USING DEPENDENCY STRUCTURE MATRIX ANALYSIS
JOHN HOLLIDAY

HDEV07: BUILDING CUSTOM APPLICATIONS (MASHUPS) ON THE SHAREPOINT PLATFORM
TODD BAGINSKI

HDEV08: BUSINESS CONNECTIVITY SERVICES (BCS) DEVELOPMENT PATTERNS
TODD BAGINSKI

HDEV09: INTEGRATING WINDOWS 7 MOBILE APPLICATIONS WITH SHAREPOINT SITES
TODD BAGINSKI

HDEV10: UPGRADING WEB PARTS FOR USE ON SHAREPOINT 2010
MAURICE PRATHER

HDEV11: BUILDING CLAIMS-AWARE APPLICATIONS AND CONTROLS
MAURICE PRATHER

HDEV12: SHAREPOINT GUIDANCE - DEVELOPING APPLICATIONS – FOUNDATION AND EXECUTION
ROBERT L. BOGUE

HDEV13: ENHANCING THE SHAREPOINT SOCIAL EXPERIENCE WITH THE SHAREPOINT 2010 SOCIAL API
MATT MCDERMOTT

HDEV14: EXPLOITING THE "HIDDEN GEMS" OF THE SHAREPOINT SOCIAL API
MATT MCDERMOTT

HDEV15: ECM FROM A DEVELOPER'S PERSPECTIVE
PAUL SWIDER

HDEV16: BUILDING APPLICATIONS WITH THE CLIENT OBJECT MODELS
SCOT HILLIER

HDEV17: ADVANCED SEARCH-BASED SOLUTIONS IN SHAREPOINT 2010
SCOT HILLIER

HDEV18: DEVELOPING RICH CLIENT SOLUTIONS WITH BUSINESS CONNECTIVITY SERVICES
SCOT HILLIER

ADMIN TRACK

HITP01: WISH I'D HAVE KNOWN THAT SOONER! SHAREPOINT INSANITY DEMYSTIFIED
DAN HOLME

HITP02: SHAREPOINT 2010 DEPLOYMENT DEMOFEST
BEN CURRY

HITP03: ARCHITECTING A SHAREPOINT SERVER 2010 FARM
BEN CURRY

HITP04: ARCHITECTURE BEHIND THE SOCIAL COMPUTING PLATFORM IN SHAREPOINT 2010
BEN CURRY

HITP05: DESIGNING GOVERNANCE: HOW INFORMATION MANAGEMENT AND SECURITY MUST DRIVE YOUR DESIGN
DAN HOLME

HITP06: A PRACTICAL JUMP START TO ADMINISTERING SHAREPOINT WITH WINDOWS POWERSHELL
DAN HOLME

HITP07: INFORMATION ARCHITECTURE AND THE MANAGED METADATA SERVICE: A TO Z
DAN HOLME

HITP08: WINDOWS POWERSHELL FOR SHAREPOINT ADMINISTRATORS AND DEVELOPERS
DON JONES

HITP09: SHAREPOINT SERVICE ARCHITECTURE DRILL-DOWN
JOEL OLESON

HITP10: UPGRADING TO SHAREPOINT 2010
JOEL OLESON

HITP11: SHAREPOINT SEARCH CHALLENGES AND TRICKS
MATTHEW MCDERMOTT

HITP12: BUILDING THE PERFECT SHAREPOINT 2010 FARM: REAL-WORLD BEST PRACTICES FROM THE FIELD
MICHAEL NOEL

HITP13: ARCHITECTING A FAULT TOLERANT AND HIGH PERFORMANCE SHAREPOINT 2010 FARM
MICHAEL NOEL

HITP14: PLANNING EXTRANET ENVIRONMENTS WITH SHAREPOINT 2010
MICHAEL NOEL

HITP15: CLAIMING TO GET FORMS-BASED AUTHENTICATION
ROBERT L. BOGUE

HITP16: PROTECT YOUR SHAREPOINT FARM FROM THE EVIL DEVELOPERS
ROBERT L. BOGUE

NO CODE SOLUTIONS TRACK

HNCS01: MANAGE YOUR EXTERNAL DATA USING BUSINESS CONNECTIVITY SERVICES ... WITHOUT CODE
ASIF REHMANI

HNCS02: USE DATA VIEWS TO GET TO YOUR DATA – BOTH INSIDE AND OUTSIDE OF SHAREPOINT
ASIF REHMANI

HNCS03: AUTOMATING BUSINESS PROCESSES USING INFOPATH 2010 FORMS WITH INTEGRATED SHAREPOINT DESIGNER 2010 WORKFLOWS
ASIF REHMANI

HNCS04: USING INFOPATH 2010 AND SHAREPOINT DESIGNER 2010 TO MANAGE SHAREPOINT LIST FORMS
ASIF REHMANI

HNCS05: PERFORMANCEPOINT SERVICES 2010: BUILDING A DASHBOARD IN 60 MINUTES OR LESS
DARRIN BISHOP

HNCS06: UNDERSTANDING POWERPIVOT AND WHAT IT BRINGS TO THE TABLE
MAURICE PRATHER.

HNCS07: SOLUTIONS WITHOUT SEMICOLONS – THE IT PROS GUIDE TO SOLUTION CREATION
ROBERT L. BOGUE

HNCS08: USING OUTLOOK AND THE SHAREPOINT WORKSPACE WITH SHAREPOINT 2010
SCOT HILLIER

HNCS09: SHAREPOINT SOLUTIONS FOR INFORMATION TECHNOLOGY PROFESSIONALS
PAUL SWIDER

ABSTRACTS ARE AVAILABLE ONLINE

SHAREPOINT AND BUSINESS TRACK

HSB01: SHAREPOINT BRANDING: CREATING A SUCCESSFUL BRANDING PROJECT MAP
CATHY DEW

HSB02: CREATING CONSISTENCY IN USER INTERFACE DESIGN WITH SHAREPOINT 2010
CATHY DEW

HSB03: DON'T JUST MIGRATE – TRANSFORM YOUR SHAREPOINT ENVIRONMENT
CHRISTIAN BUCKLEY

HSB04: SHAREPOINT'S SOCIAL COMPUTING SCORECARD
CHRISTIAN BUCKLEY

SQL SERVER CONNECTIONS SESSIONS

SDV304: SPATIAL DATA "STRETCHES OUT" IN SQL SERVER DENALI
ROBERT BEAUCHEMIN

SDV306: FULL TEXT SEARCH IN SQL SERVER 2008 AND DENALI
ROBERT BEAUCHEMIN

SDB414: EXPANDING THE SCOPE AND EASE OF USE OF QUERY PLAN GUIDES IN SQL SERVER 2008
ROBERT BEAUCHEMIN

SDV305: GETTING SQL SERVICE BROKER UP AND RUNNING
DENNY CHERRY

SDB206: SQL SERVER CLUSTERING 101
DENNY CHERRY

SDB215: EXPLORING THE DAC AND EVERYONE'S FAVORITE FEATURE, THE DACPAC
DENNY CHERRY

SDB202: REMOTE BLOB STORAGE: THE QUESTIONS AND THE ANSWERS
VICTOR ISAKOV

SDB311: TROUBLESHOOTING PARALLELISM PROBLEMS IN SQL SERVER
VICTOR ISAKOV

SDB203: WHAT EVERY DBA SHOULD KNOW ABOUT SHAREPOINT 2010
VICTOR ISAKOV

SBI102: BUILD YOUR FIRST SSIS PACKAGE
ANDY LEONARD

SDV201: DATABASE DESIGN FOR DEVELOPERS
ANDY LEONARD

SBI204: INTRODUCTION TO INCREMENTAL LOADS
ANDY LEONARD

SBI201: CREATING REPORT SUBSCRIPTIONS IN MICROSOFT SQL SERVER 2008 REPORTING SERVICES
PAUL LITWIN

SBI303: PROGRAMMING SQL SERVER 2008 REPORTING SERVICES
PAUL LITWIN

SDV207: TUNING T-SQL STEP BY STEP
BRENT OZAR

Sessions and speakers are subject to change.
Check the Web site for details.

HSB05: TRUST ME I AM A DEVELOPER: THINGS AN ADMIN SHOULD KNOW ABOUT DEVELOPING ON SHAREPOINT
DARRIN BISHOP

HSB06: SHAREPOINT AS A PLATFORM FOR BUSINESS APPLICATIONS
OWEN ALLEN

HSB07: HORIZONTAL AND VERTICAL BUSINESS SOLUTIONS FOR SHAREPOINT 2010
OWEN ALLEN

SDB301: BLITZ! SQL SERVER TAKEOVERS
BRENT OZAR

SDB207: CONSOLIDATION, CLUSTERING, AND VIRTUALIZATION: CHOOSING WISELY
BRENT OZAR

SDB409: "DUDE, WHERE IS MY MEMORY?" UNDERSTANDING MICROSOFT SQL SERVER MEMORY USAGE AND MANAGEMENT
MACIEJ PILECKI

SDV308: SQL SERVER USER-DEFINED FUNCTIONS - THE GOOD, THE BAD, THE UGLY
MACIEJ PILECKI

SDB304: TROUBLESHOOTING DEADLOCKS IN SQL SERVER
MACIEJ PILECKI

SDB208: MORE DBA MYTHBUSTERS
PAUL S. RANDAL

SDB310: INDEX FRAGMENTATION: THE HIDDEN MENACE
PAUL S. RANDAL

SDB412: UNDOCUMENTED TOOLS AND TRACE FLAGS
PAUL S. RANDAL

SDB213: FOLLOW THE RABBIT: WRAP-UP Q&A
PAUL S. RANDAL

SDB305: VLDB: RECOVERING FROM ISOLATED DISASTERS
KIMBERLY L. TRIPP

SDV302: INDEX INTERNALS: WHAT YOU REALLY NEED TO KNOW!
KIMBERLY L. TRIPP

SDV303: OPTIMIZING PROCEDURAL CODE
KIMBERLY L. TRIPP

ABSTRACTS ARE AVAILABLE ONLINE

Register Today!
Call 800-438-6720
www.WinConnections.com

PRE-CONFERENCE SESSIONS
MARCH 27, 2011

9AM - 12PM

WPR01: AUTOMATING ACTIVE DIRECTORY ADMINISTRATION
MARK MINASI

1PM - 4PM

WPR02: GROUP POLICY FUNDAMENTALS, SECURITY, AND CONTROL
JEREMY MOSKOWITZ

9AM - 4PM

WPR03: WINDOWS 7 DEPLOYMENT MASTER CLASS
RHONDA LAYFIELD

HPR301: SHAREPOINT 2010 PROFESSIONAL DEVELOPMENT
ROBERT L. BOGUE & ERIC SCHUPPS

HPR302: DAN HOLME'S SHAREPOINT COLLABORATION MASTERCLASS
DAN HOLME

SPR302: DAY OF SCRIPTING: PLUMBING THE DEPTHS OF SQL SERVER / POWERSHELL INTEGRATION
BOB BEAUCHEMIN

SPR301: THE BUILDING BLOCKS OF A HEALTHY AND AVAILABLE SQL SERVER
KIMBERLY L. TRIPP & PAUL S. RANDAL

The final list of Exchange workshops were not confirmed at press time. Check online for the workshop descriptions.

POST-CONFERENCE SESSIONS
MARCH 31, 2011

9AM - 12PM

WPS01: MIGRATING AND RESTRUCTURING YOUR AD
J. PETER BRUZZESE

9AM - 4PM

HPS301: BUSINESS CONNECTIVITY DEEP DIVE
SCOT HILLIER & TODD BAGINSKI

HPS302: ORGANIZING INFORMATION IN SHAREPOINT SERVER 2010
BILL ENGLISH

SPS301: INDEXING STRATEGIES AND ANALYSIS
KIMBERLY L. TRIPP

SPS302: VIRTUALIZATION AND SAN BASICS FOR DBAS
BRENT OZAR

The final list of Exchange workshops were not confirmed at press time. Check online for the workshop descriptions.

ABSTRACTS ARE AVAILABLE ONLINE

Featured Speakers

J. PETER BRUZZESE
CLIP TRAINING

MIKE DANSEGLIO
CONCENTRATED TECHNOLOGY

SEAN DEUBY
WINDOWS IT PRO

DEVIN L. GANGER
CONSULTANT/AUTHOR

LAURA HUNTER
MICROSOFT

DON JONES
CONCENTRATED TECHNOLOGY

KEVIN LAAHS
HP

RHONDA LAYFIELD
CONSULTANT/TRAINER

JIM MCBEE
ITHICOS SOLUTIONS

KIERAN MCCORY
HP

MARK MINASI
MINASI RESEARCH AND DEVELOPMENT

JEREMY MOSKOWITZ
MOSKOWITZ, INC.

GREG SHIELDS
CONCENTRATED TECHNOLOGY

MICHAEL B. SMITH
THE ESSENTIAL EXCHANGE

PETER O'DOWD
DATACOM/WADEWARE

ALAN SUGANO
ADS CONSULTING GROUP

"THE CONVERSATION BEGINS HERE"

Microsoft® Exchange CONNECTIONS

WINDOWS CONNECTIONS

UNIFIED COMMUNICATIONS CONNECTIONS

SharePoint CONNECTIONS

SQL Server CONNECTIONS

Microsoft® Visual Studio CONNECTIONS

Microsoft® ASP.net CONNECTIONS

Microsoft® Silverlight CONNECTIONS

BONUS:
Mobile Apps Track

MARCH 27-30, 2011 • ORLANDO, FL
GRANDE LAKES JW MARRIOTT RESORT HOTEL

Book NOW to get a special rate (a limited number of rooms at this rate, so reserve today).

CHECK WEB SITE FOR DESCRIPTIONS OF SESSIONS AND WORKSHOPS

www.DevConnections.com • 800.438.6720 • 203.400.6121 • Register Early!

Database Maintenance for SharePoint

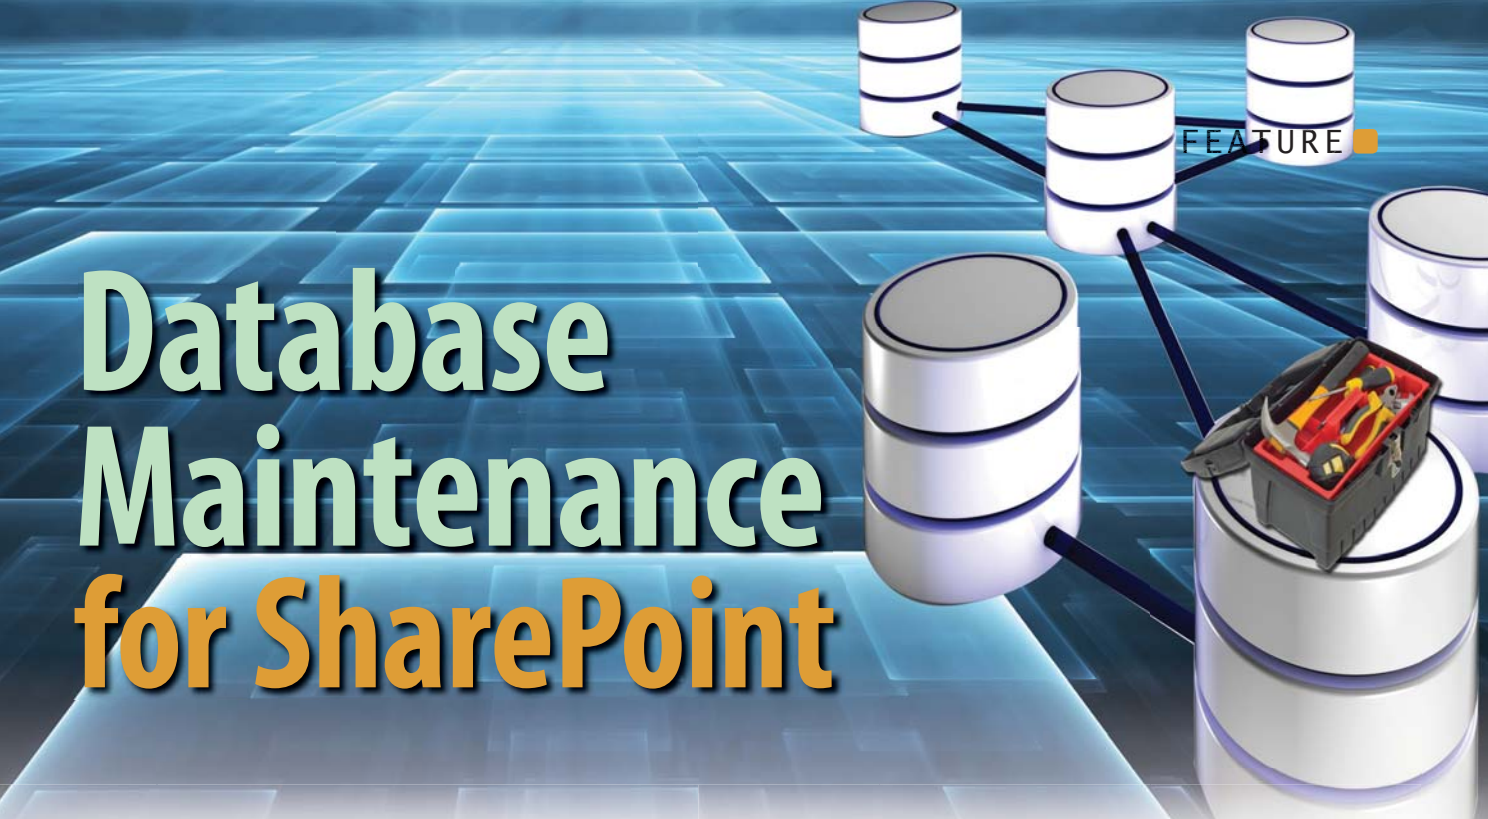

You can easily keep SharePoint's databases in optimum condition

by Matt Ranlett and Brendon Schwartz

SharePoint administrators are defenders of order in any SharePoint installation. The integrity of SharePoint data comes from proper organization of sites, site collections, content types, and metadata tags. As owners of the configuration and functionality of SharePoint, administrators share the responsibility in the overall performance and stability of the SharePoint platform. While ensuring the optimum performance of a SharePoint farm takes more than a well-oiled set of databases, it certainly helps to keep the lowest-level components operating as smoothly as possible.

Fortunately, you can keep SharePoint's databases in optimum condition with standard maintenance tasks. It's important for SharePoint administrators to understand what those maintenance tasks are and to know the best practices for how and when to run them. SharePoint 2010 ships with Microsoft SQL Server Maintenance Plan Wizard, which helps SharePoint administrators to automate these tedious maintenance tasks, but it's important to understand exactly what needs to be done.

Databases

SharePoint 2007 was a simple affair when it came to understanding and managing the databases required for healthy operation of the application, with only a content database, a configuration database, and perhaps a search database or two. SharePoint 2010 is a different animal. With the introduction of service applications and the large set of core databases, one SharePoint 2010 installation might have as many as 20 databases, each of which needs to be backed up and maintained. Therefore, it's important to know which databases are core to the system or required for the proper functioning of various SharePoint service applications (Figure 1 and Figure 2).

Before diving too deeply into SharePoint databases, it's important to understand which versions of SQL Server are supported by SharePoint 2010 installations. SharePoint 2010 supports 64-bit versions of SQL Server 2008 R2, SQL Server 2008, and SQL Server 2005. We recommend you use SQL Server 2008 R2 because it has the widest set of functionality and features for SharePoint. For example, Remote Blog Storage is available only on SQL Server 2008 Enterprise. For more information about the different SQL Server versions, see the whitepaper "SQL Server 2008 R2 and SharePoint 2010: Better Together" at technet.microsoft.com/en-us/library/cc990273.aspx.

Be sure that any SQL Server setup is up to date with the latest patches and hotfixes per the recommendations in the TechNet article "Hardware and Software Requirements (SharePoint Server 2010)" at technet.microsoft.com/en-us/library/cc262485.aspx. This page is kept up to date as new cumulative updates and patches are released. Be sure to visit this page at least once a month to ensure compliance. At the time of this writing, the page had been updated three times, most recently in July 2010.

Configuration Databases

Name	Size and Growth	Read/Write Characteristics	Scaling Method	Recovery Model	Backup Mechanisms	Mirror Within Farm	Mirror/Log Ship to Another Farm
SharePoint_ AdminContent_ GUID	Small	Varies	Only scales up; one database per farm	Full	SharePoint backup, SQL Server, Data Protection Manager (DPM)	Yes	No
SharePoint_Config	Small with large log files	Read-intensive	Only scales up; one database per farm	Full	SharePoint backup, SQL Server, DPM	Yes	No

Content Databases

Name	Size and Growth	Read/Write Characteristics	Scaling Method	Recovery Model	Backup Mechanisms	Mirror Within Farm	Mirror/Log Ship to Another Farm
WSS_Content	Recommended to limit size to < 200GB*; up to 1TB supported in some scenarios	Varies	Can scale up or split additional site collections to new databases	Full	SharePoint backup, SQL Server, DPM	Yes	Yes

Health Databases

Name	Size and Growth	Read/Write Characteristics	Scaling Method	Recovery Model	Backup Mechanisms	Mirror Within Farm	Mirror/Log Ship to Another Farm
WSS_ UsageApplication	Extra large	Write-intensive	Only scales up; one database per farm	Simple	SharePoint backup, SQL Server, DPM	Yes, but why bother	Yes, but why bother
WSS_Logging							

Application Databases

Name	Size and Growth	Read/Write Characteristics	Scaling Method	Recovery Model	Backup Mechanisms	Mirror Within Farm	Mirror/Log Ship to Another Farm
Application_Registry_ Service_DB_GUID	Small	Read-intensive	Only scales up; one database per farm	Full	SharePoint backup, SQL Server, DPM	Yes	No
Bdc_Service_DB_ GUID	Small	Read-intensive	Only scales up; one database per farm	Full	SharePoint backup, SQL Server, DPM	Yes	No

Figure 1: SharePoint 2010 Foundation Databases

DBA Created vs. Auto-Provisioned Databases

With few exceptions, SharePoint 2010 will create all the databases required for the healthy operation of a SharePoint environment. Given the ease of auto-provisioning, when should a DBA insist on pre-creating SharePoint databases? In production environments, certain circumstances dictate a more prudent approach than auto-provisioning, including situations where you need

- guaranteed control over database names (no GUIDs in the database names)
- guaranteed control over database sizing
- procedural separation of control over application and data environments

If you decide to create databases by hand rather than through an auto-provisioning technique, you must use the appropriate PowerShell cmdlets to create the databases and register them with SharePoint. For example, you can use the following PowerShell command to create a new configuration database:

```
New-SPConfigurationDatabase
  -DatabaseName "SharePointConfigDB1"
  -DatabaseServer "SQL-01"
  -Passphrase (ConvertTo-SecureString
  "MyPassword" -AsPlainText -force)
  -FarmCredentials (Get-Credential)
```

You can find a complete walkthrough of how a DBA might create various content and configuration databases in the TechNet article "Deploy by using DBA-created databases (SharePoint Server 2010)" at technet.microsoft.com/en-us/library/cc262869.aspx.

Data Integrity

Nothing can ruin a business-critical content repository's reputation faster than data corruption. As the DBA or administrator responsible for the quality of the platform, it's important that you understand database corruption and how you can correct it.

Protecting data is tricky, especially if errant power spikes and sags cause your

- Managed Metadata Service_GUID
- PerformancePoint Service Application_GUID
- Search_Service_Application_CrawlStoreDB_GUID
- Search_Service_Application_DB
- Search_Service_Application_PropertyStoreDB_GUID
- Secure_Store_Service_DB_GUID
- StateService_GUID
- User Profile Service Application_ProfileDB_GUID
- User Profile Service Application_SocialDB_GUID
- User Profile Service Application_SyncDB_GUID
- WebAnalyticsServiceApplication_ReportingDB_GUID
- WebAnalyticsServiceApplication_StagingDB_GUID
- WordAutomationServices_GUID

Figure 2: SharePoint 2010 Server Databases

SQL Server's I/O subsystem to fail when writing to disk. Corruption is frequently the result, and without the proper checks, that corruption will lurk in your database undetected until you need quality data. Corruption in a database might occur when the disk that holds a log file or data file has been altered. This type of physical corruption tends to affect sectors on a disk due to problems in the I/O subsystem, such as the physical network hardware and disk drives themselves. Physical corruption, which is what the built-in SQL function DBCC CHECKDB looks for and reports, is usually caused by physical hardware problems.

Logical corruption is caused by data being altered in some unanticipated way that severs a data relationship. This type of corruption usually is caused by an application error or human error that causes data problems but doesn't affect the physical structure of the database. Bugs can also cause this type of corruption, which you can learn about in the Microsoft article "When you use a file handle for FileStream access in a SQL Server 2008 transaction, the transaction may randomly fail to commit in Windows Server 2003 or in Windows XP Professional x64 Edition" (support.microsoft.com/?kbid=955280).

Best practice dictates that you should run a DBCC CHECKDB command at least as often as you run a full backup. DBCC CHECKDB presents a report of errors that you can investigate further. It's important to note that DBCC CHECKDB doesn't check for logical corruption. But the command can cause logical corruption when the REPAIR_ALLOW_DATA_LOSS option is used because the option doesn't take any constraints into consideration when repairing physical corruption issues.

When a DBCC CHECKDB command returns with error messages, the appropriate solution is to turn to your database backups. However, this requires regular backups that aren't corrupted. As mentioned earlier, DBCC CHECKDB can potentially introduce errors if it's used to fix corruption. Without backups, there's no way to get sanitized data from a corrupt database.

Speed

SharePoint is a complex application that relies on many different frameworks, components, and server applications to function and perform properly. Improper maintenance at any of these levels can contribute to problems and perhaps downtime. Fortunately, configuring SQL Server for optimum performance isn't complicated. Performance tuning SQL Server largely involves a couple of configuration settings, proper placement of data and log files, and the occasional rebuild of table indices.

First is the proper selection of SQL Server hardware. As identified in Microsoft's "SQL Server 2008 R2 and SharePoint 2010: Better Together" whitepaper (technet.microsoft.com/en-us/library/cc990273.aspx), SharePoint 2010 requires a 64-bit SQL Server installation because 32-bit hardware and software are no longer supported. Basic RAID disk configuration, SAN slice allocations, and local storage requirements also play a major role in determining the performance of SQL Server because several recommendations require the physical separation of data files and log files onto different spindles. Finally, memory requirements for the SQL Server system start at 16GB and move up from there. Microsoft provides some SharePoint-specific SQL Server sizing recommendations in the TechNet article "Storage and SQL Server capacity planning and configuration (SharePoint Server 2010)" at technet.microsoft.com/en-us/library/cc298801.aspx.

Managing Database Files

When it comes to managing SQL Server data files, best practices dictate that data and log files reside on their own physical spindle. This is more than just disk volume separation; the recommendation is to place log files and data files on different disks and to ensure that no other application uses those disks. This setup minimizes overall write access to the disks, and lessens the opportunities for file fragmentation.

Another recommendation is to try to pre-create database and log files of the appropriate size ahead of time, rather than allowing the files to auto-grow by small increments. The reason for this configuration is that an auto-grow operation can take time and slow down write-intensive environments such as SharePoint collaboration farms. This doesn't mean that DBAs should disable auto-grow on SharePoint databases and logs, but you shouldn't rely on this capability for the initial sizing of the databases. Note that the behavior of auto-grow for SQL Server 2005 and SQL Server 2008 has changed. In earlier versions of SQL Server, log files were initialized and zero-filled when auto-grown, which is a large part of why the operation was slow. Proper configuration of SQL Server 2005 and SQL Server 2008's instant file initialization option allows for the elimination of that zero-filling initializing step.

Measuring and Reducing Fragmentation

Data fragmentation inside of SQL Server usually can be explained by normal data manipulation, including inserts, updates, and deletes. The basic symptom of data fragmentation is an increasing volume of free space corresponding to data volumes. This is because SQL Server stores data as a database page, which contains header details, record details, and an index. However, the database page is also configured by the SQL Server fill factor to have a minimum size. Records smaller than the fill factor result in a lot of empty space in a table. SharePoint systems rely on GUID-based index keys and therefore exacerbate the problem by inserting data randomly throughout the range of records.

To correct this problem, best practices suggest either changing the index schema or rebuilding the index to compact and defragment the data. Because SharePoint can't have its index schema changed, the only option is to rebuild the index for tables or indexes that require it. The temptation is to periodically rebuild all indexes using an automated maintenance plan, but this tends to require significant amounts of available free space and is a poor choice for large systems. A more appropriate choice would be to use the DMV SYS.DM_DB_INDEX_PHYSICAL_STATS to find the indexes that are most fragmented. MSDN provides sample code at msdn.microsoft.com/library/ms188917.

Keeping Statistics Up to Date

When SQL Server executes a query, it does so along a calculated and compiled execution plan. The execution plan is created by SQL Server's Query Processor, and it defines which tables and indexes to use to achieve the best possible query performance. The metrics for determining query performance

are derived from statistics that help SQL Server understand how data is distributed inside a table or index. These statistics are generated by a variety of read operations, including full and sample data scans. If these statistics are out of date due to fragmented indexes, the execution plan will not be as efficient as it could otherwise be.

Statistics are usually kept up to date automatically, but certain maintenance operations or intensive data manipulation can cause them to become out of date. You can force SQL Server to update the statistics via the built-in SP_UPDATESTATS function. Running SP_UPDATESTATS after an index rebuild isn't recommended because this changes the previous sampling level (determined automatically) and might result in less than optimal results. Per the August 2008 TechNet article "Top Tips for Effective Database Maintenance" (technet.microsoft.com/en-us/magazine/2008.08.database.aspx), the recommended database maintenance plan performs the following steps:

- Analyze indexes and determine which indexes to operate on and how to do the fragmentation removal.
- For all indexes that weren't rebuilt, update the statistics.
- Update statistics for all the non-indexed columns.

Data Sizing

Right-sizing your data tier is a critical part of any SharePoint implementation because it has direct impact on the cost, performance, and scalability of the entire application. To understand how much data your SharePoint environment must be able to hold, you need to consider the total number of documents, the number of versions of each document, and the average size of each document. Additionally, you must consider the number of list items that will be stored in the application. Microsoft has provided the formula below for data sizing. You can read more in the TechNet article "Storage and SQL Server capacity planning and configuration (SharePoint Server 2010)" (technet.microsoft.com/en-us/library/cc298801.aspx).

```
Database size = ((# Documents × #
  Versions) × Avg Size) + (10 KB
  ×(# List items + (# Versions × #
  Documents)))
```

Although this simple calculation will give you a rough idea of the required storage capacity, you do several things to affect the resulting number. As previously mentioned, database content is stored in database pages, which are kept to a uniform size by the configured fill factor. Adjusting the fill factor can affect fragmentation and overall size on disk. It's possible to actually shrink a database, although Microsoft is issuing some fairly strong warnings these days about wonton use of the DBCC_SHRINKDATABASE function. Finally, if allocating all data storage responsibilities to SQL Server entails too much cost and overhead, you can configure SharePoint to look to the file system for storage of large binary objects or BLOBs.

Set the Fill Factor for a Server

By changing the default configuration for FILLFACTOR, the DBA can reduce fragmentation and page splits (a symptom of fragmentation that affects performance), but the side effect is that this takes more database space because the database pages are larger. The database fill factor defines how much free space is required on a database page during an index rebuild before moving to a new database page. During regular operation of the database, new content can be inserted into this free space without requiring a clustered index to adjust large amounts of data. Kimberly Tripp offers a great deal of information about this important attribute in her blog series on database maintenance best practices at www.sqlskills.com/BLOGS/KIMBERLY/post/Database-Maintenance-Best-Practices-Part-II-e28093-the-most-important-setting-FILLFACTOR.aspx.

Shrinking Databases

All the supported versions of SQL Server for SharePoint have the ability to shrink data files, which recovers disk space by removing unused data. None of the databases in SharePoint are set to automatically shrink the data files. The strong recommendation from Microsoft and knowledgeable SQL Server MVPs is not to perform auto-shrink of the database or to configure a maintenance plan that does it automatically on a database. The reason is that the shrink ignores the fill factor setting and causes all the indexes to become fragmented. Then, when you run a rebuild indexes command, the database grows back to its original size. Instead of relying on SQL Server's automated DBCC_SHRINKDATABASE commands, it's safer to partition content databases or to remove data from existing databases. The following list shows the activities you can perform to create free space in a SharePoint environment:

- Use STSADM MERGCONTENTSDB
- Delete documents
- Delete libraries
- Delete lists
- Delete list items
- Delete sites

Remote BLOB Storage

To free up critical resources such as the file system, you can use the new Remote BLOB Storage (RBS) mechanism available with SQL Server 2008. Although this sounds like a great way to move large data items (e.g., image files, streaming video or sound clips) out of your database, you need to evaluate the advantages and disadvantages first. If the files are not large, and you have many small BLOBs, you can see a decrease in performance on the server.

So make sure to evaluate your content to determine whether you need to implement the RBS. The current recommendation is that your content database should be larger than 500GB and the BLOB data files larger than 256KB. RBS will provide the most performance increase on systems that have large to extremely large files that aren't frequently accessed. Adding RBS to your write-intensive SharePoint collaboration implementation could actually make the user experience slower.

Database Maintenance Plans

Almost all tasks that you can perform with SQL Server can be automated. It's crucial to any SharePoint implementation to have an automated plan that will help maintain your site. Keep track of when these automated tasks run because you might have to notify users or plan for the system to run a little slower during this time to keep everything in tip-top shape.

A database maintenance plan is like routine upkeep on your car. You must do maintenance on a regular schedule

to have your SQL Server deployment perform at an optimum level and to help keep your website running at peak performance.

The great part about the SQL Server Maintenance Plan is that you can use the Maintenance Plan Wizard to set it up. The wizard gives you the ability to add items such as database backups and transaction logs, to update database statistics, and to manage data such as indexes.

Based on the August 2008 TechNet magazine article "Top Tips for Effective Database Maintenance" (technet.microsoft.com/en-us/magazine/2008.08.database.aspx), here's a checklist of tasks that you should include in your maintenance plan:

- Remove excessive transaction log file fragmentation by ensuring the appropriate recovery model and backup schedule.
- Turn off any scheduled shrink operations to reduce the risk of unnecessary index fragmentation.
- Set auto-growth correctly by using a predetermined set file size rather than a percentage. Follow this up by periodically examining database sizes and determining whether manual database growth is necessary to ensure optimum performance.
- Turn on instant file initialization such that database auto-growth is an instantaneous operation rather than a slow operation that requires zero-filling new space.
- Put a regular process in place to detect and remove index fragmentation.
- Turn on AUTO_CREATE_STATISTICS and AUTO_UPDATE_STATISTICS, and have a regular process in place to update statistics.
- Turn on page checksums.
- Have a regular process to run DBCC CHECKDB.

Any SharePoint administrator should make sure to spend a little time creating at least a basic maintenance plan to keep his or her system running well. The time spent setting it up will pay off with performance, speed, and backups, if ever needed.

InstantDoc ID 126012

Matt Ranlett
(mranlett@devcow.com), a SharePoint Server MVP, works as a solution architect with Slalom Consulting. Matt focuses on strategic technical implementations for enterprise clients in the Atlanta area and has been an author and editor on six different books. Visit his blog at blogs.sharepointguys.com/matt.

Brendon Schwartz
(bschwartz@devcowsoftware.com) is a SharePoint MVP, working on cutting-edge SharePoint projects for the past 7 years. He was an author of 2 SharePoint books and technical editor for *Beginning SharePoint 2010 Development* (Wrox). Visit his blog at blogs.sharepointguys.com/brendon.

Enable Multi-User Video Conferencing

VideoPort has released **VideoPort Online 6.0.1,** a software server designed to set up a multi-user video conferencing system for corporate work groups. Several features include: virtual business meetings, file sharing and collaboration, online presentations, text chat, and remote desktop sharing. New to this version are multipoint video conferencing, video lecture, and video meeting modes. Video and audio conference sessions can be recorded in a specified format. Also, chat conversations can be logged and stored for future reference. For more information about this product, visit www.video-port.com.

Power Admin's PA Server Monitor 4.0

Power Admin has released **PA Server Monitor 4.0,** a server monitoring tool. PA Server Monitor 4.0 can monitor many types of server and network resources including: event logs; CPU, memory, and network usage; disk space; running services; log files; running processes; and directory quotas. The new version of PA Server Monitor supports satellite monitoring of remote client sites across the Internet. In addition, users can work from a remote desktop. A number of status reports are available, including group summaries, uptime reports, and history statistics. Reports can be viewed via HTTP, password protected, and delivered via email. For more information, visit www.poweradmin.com.

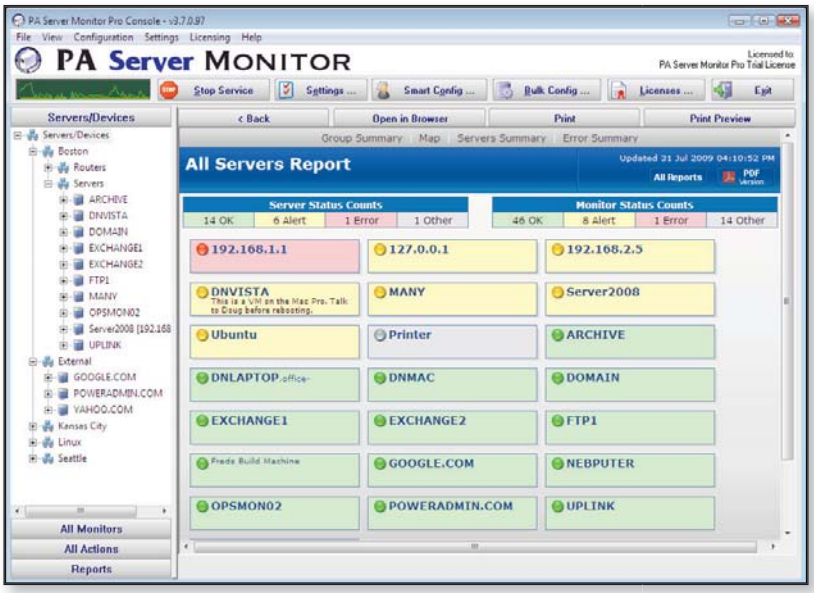

Best of Connections 2010 Awards

In early November, we handed out our Best of Connections 2010 awards at the annual Fall Connections conference. The winners were chosen from more than 80 nominated products in six categories. The field was narrowed to three finalists which were interviewed on the show floor to determine the winners. Congratulations to the winners and finalists!

Best Windows Product

- Winner: STEALTHbits Technologies StealthAUDIT Management Platform
- Finalists: ScriptLogic Enterprise Security Reporter and Imanami GroupID

Best SQL Server Product

- Winner: SQLSentry Performance Advisor
- Finalists: Idera SQL diagnostic manager and Red Gate Software SQL Source Control

Best SharePoint Product

- Winner: Idera SharePoint diagnostic manager
- Finalists: AvePoint DocAve Software Platform and Axceler Davinci Migrator for SharePoint 2010

Best Visual Studio Product

- Winner: Infragistics NetAdvantage for .NET Ultimate
- Finalists: ComponentOne Studio Enterprise and GrapeCity ActiveReports 6

Best ASP.NET/Silverlight Product

- Winner: Aptimize Website Accelerator
- Finalists: ScaleOut Software ScaleOut StateServer and ComponentOne Studio for Silverlight

Best Exchange Server/Unified Communications Product

- Winner: Mimecast Unified Email Management
- Finalists: ENow Mailscape and Sherpa Software PST Backup Attender

BullGuard Releases Internet Security Suite 10

BullGuard has released **Internet Security Suite 10,** which provides users with protection from viruses and other online threats. New features include behavior detection; safe browsing that flags any registered or unsafe websites that come up in search results; and inspection which locates vulnerable, out-dated software before hackers exploit it. Additional features include antivirus, anti-spyware, anti-phishing, firewall protection, spam filter, backup, social media protection, password protection, and 24/7 live support. To learn more about this product, visit www.bullguard.com.

Ciphertex Debuts Secure Network Storage Appliance

Ciphertex Data Security has released **CX-5000NAS AES 256 Encryption Secure**

Network Storage Appliance, a solution for secure network storage in virtualized environments. The CX-5000NAS appliance can house up to five 3.5 inch or 2.5 inch SATA hard disk drives and incorporates an Intel Atom 1.8GHZ dual core processor with 1GB memory. The CX-5000NAS supports file sharing across Windows, Mac, Linux, and UNIX platforms. In addition, business applications such as file server, FTP server, printer server, web server, and Windows Active Directory support are provided. To learn more about this product, visit www.ciphertex.com.

Colasoft Releases Capsa Network Analyzer 7.3.1

Colasoft has released **Capsa Network Analyzer 7.3.1,** a network analyzer that detects, isolates, and resolves network problems. New features include a tab management panel of the main view and a data storage option on the Start Page for packet and log save settings. Also, there are new views in the Security Analysis Profile that analyzes Dos attack, APR attack, and worm activities. Capsa 7.3.1 runs under Windows XP, Windows 2003, Windows Vista, and Windows 7. To learn more, visit www.colasoft.com.

SharePoint Solutions Releases Alert Manager

SharePoint Solutions has released **Alert Manager for SharePoint Server 2010** and **SharePoint Foundation 2010,** an add-on that allows users to customize and manage SharePoint alerts. New features include: the ability to manage alerts across an entire site collection from one central location; subscribe other SharePoint users to alerts on any list, document library, or document; and customize alert templates with images and branding. To learn more about this product, visit www.sharepointsolutions.com.

Paul's Picks

www.winsupersite.com

SUMMARIES of in-depth product reviews on Paul Thurrott's SuperSite for Windows

Internet Explorer 9

PROS: Desktop integration features; standards compliance; hardware-accelerated performance.

CONS: Best features available only in Windows 7

RATING: ◆◆◆◆◆

RECOMMENDATION: Internet Explorer 9 is one of the most important products to come out of Redmond in years—you might say it's even more important than Windows 7. It answers all the key complaints about IE, including performance and standards compliance, and it also makes the competition look silly. Most interesting, IE 9 also includes a slew of useful new end-user features. Key among these improvements are desktop integration features—available only in Windows 7—that make websites look and work much like traditional Windows applications, with their own jump lists, notifications, and other familiar features. IE 9 is just the latest in a suddenly long list of Microsoft successes, but it's the most interesting, too, given the world's transition to web-based apps. Suddenly, the Windows desktop is important again, and the reason, curiously, is IE 9.

CONTACT: Microsoft • www.microsoft.com

DISCUSSION: www.winsupersite.com/live/ie9_beta.asp

Kinect for Xbox 360

PROS: Takes motion-sensing mainstream; enables new interaction paradigms; not just for the Xbox 360

CONS: Performance is a bit slow; facial recognition could be better

RATING: ◆◆◆◆◇

RECOMMENDATION: On the surface, Microsoft's Kinect motion sensor add-on for the Xbox 360 is an interesting development in the video game world and a way for Microsoft to simultaneously address the dominance of Nintendo's Wii while extending the life of its own console. And on that note, it succeeds. Yet the Kinect also points to deeper computer-human interface changes coming down the road, and the implications will be seen in future desktop versions of Windows as well as other products, resulting in a new, more pervasive generation of computing devices. Oddly, the Kinect's biggest gains are in voice command, which has been available on PCs and mobile devices for years but never as integrated and well-done as is the case with this console. It's a quiet revolution for now, but make no mistake: Kinect is a revolution, and one that will make an impact on far more than just games.

CONTACT: Microsoft • www.microsoft.com

DISCUSSION: www.winsupersite.com/xbox/kinect.asp

InstantDoc ID 128837

REVIEW

Viewfinity Systems Management Suite

A common problem in IT is how to manage systems, regardless of physical location, using a common tool set. This review focuses on **Viewfinity Systems Management Suite**, a set of mostly proprietary management tools for managing Windows endpoints from the cloud without installing on-premises server software.

Using a Windows XP machine, I started by signing up for a fully featured online Viewfinity trial account. After downloading and installing an .msi file, I logged back on to the website, and I was presented with a clean and straightforward UI that listed the main menu options: Computers, Policies, Reports, Deploy Agents, Software Distribution, and Patch Management. The Computers option provides the functionality to inventory a computer's hardware and software configuration. It also lets you perform common administrative functions, such as checking event logs, installing software, and remotely connecting to another desktop.

There are multiple options for locating computers for agent deployment. I chose Active Directory domain membership to roll out the agent to several Windows 7 and Windows Server 2003 systems—but only after a few false starts. Turning off the default Windows firewall gave me the best results with agent deployment, although opening just a few firewall ports will normally suffice. For remote users, an "email me the install link" option is available.

Only a small amount of system data is actually stored in the cloud. When I accessed a Windows event log through Viewfinity, the system was slow to retrieve information over the Internet. According to the system documentation, data is stored on the managed PCs due to security concerns. However, given this focus on security, it was perplexing that Viewfinity didn't use two-factor authentication for logon to the management website. I expected bank-like security, since you can remotely log on to managed machines from the Viewfinity website.

Viewfinity has a range of features that let you manage a PC without connecting via Remote Desktop, and I found all of them to be useful. I used the system-restore feature several times. Without logging on to the source PC, I was able to roll back the machine to a restore point after misconfiguring a system. Because Viewfinity includes excellent status and software deployment tools, I used the bare-bones remote control solution less often than I normally would with other PC management packages. That aside, the remote control options allowed me to log on to the target PCs, exchange files, and chat with users.

To keep Windows patches current, Viewfinity handles Windows updates via a semi-customized Windows Software Update Services (WSUS) program. I configured the WSUS system with just a few clicks through a well-designed wizard, and since Viewfinity automatically configures the auto-update options for Windows Updates, I didn't need to configure any Active Directory Group Policy settings. After approving patches via the console, I checked the C:\Windows\WindowsUpdate.log file on a target PC and confirmed that the Viewfinity server was the patch approval source and that updates were downloaded from the Microsoft website. Comprehensive reporting showed all patch successes and failures.

Software deployment is a simple two-step process. First, upload the software to Viewfinity's cloud, or point to a Windows share. Then, configure software deployment targets based on several options, including groups or single computers. To test an .exe file distribution, I uploaded the Cute PDF writer application to Viewfinity's application cloud, selected a PC group for deployment, and pushed the software. For .exe files, I found that typically the user has to click through the installation routine. So, before you push any software, it's critical to test the deployment locally.

MSI package pushes went smoother. I uploaded a Skype .msi file and a Snagit .msi file to the Viewfinity cloud (I could have specified a network share), and Viewfinity built the silent install script. Just as encouraging, the install succeeded. There are also helpful options for detecting whether the program already exists. I easily configured Viewfinity to check for the existence of a directory before installing an application. Comprehensive status information categories—such as *scheduled*, *in process*, *deployed*, *success*, and *failure*—ensured that I always knew the deployment status of a package. Additionally, Viewfinity's report engine made it easy to configure the reporting of almost any data tracked in the system. This report engine appears to be based on Microsoft Reporting Services. Email is also a handy delivery option for scheduled reports.

After using Viewfinity for more than a month, I found that the system was easy to use. And, best of all, it allowed me to resolve problems without logging on to user machines because a multitude of data is tracked and consolidated in the Viewfinity web interface. Managing machines inside and outside the firewall has never been easier!

InstantDoc ID 129052

Viewfinity Systems Management Suite

PROS: Intuitive and efficient UI; excellent status information and reporting; good software deployment tools

CONS: Logon security too basic

PRICE: $30 per endpoint for an annual subscription with software support and maintenance; tiered volume discounts available

RATING: ◆◆◆◆◆

RECOMMENDATION: Viewfinity has an impressive range of proprietary features for managing Windows systems. If you've been struggling to manage systems beyond the firewall, Viewfinity is definitely worth a try.

CONTACT: Viewfinity • 781.522.7474 • www.viewfinity.com

Tony Bieda | tonybieda@yahoo.com

Network Inventory Advisor

IT professionals need to know what's going on in the networks they manage, and an important part of that is having a detailed inventory of the computers and devices that are connected. This arduous task is often borne out of necessity: An audit is being conducted or a software upgrade is planned, and you need to know, for example, whether the computers on your network are ready to support the upgrade without requiring additional RAM.

If you're in the market for a software tool to collect such inventory details, you'll find that there's a plethora of software products available. On the surface, they all do the same basic thing: They scan your network, attempt to connect to any computers and devices they find, and pull down inventory details from anything they were able to connect to.

Where these products differentiate themselves is in the inventory details that they capture and the reporting that they produce, which were my primary areas of interest when reviewing **ClearApps Network Inventory Advisor.**

ClearApps offers a 15-day, fully featured downloadable trial version of its product. You don't have to sign up for a follow-up sales call or register for an account on their website to access the trial, which is always a very welcome touch and eliminates one of the main barriers that prevents folks from giving products like this a shot.

After a quick installation process, you're asked to provide some simple details about your network, including the number of computers (or nodes) that you have, whether you're interested in looking for non-Windows nodes (including SNMP devices and Macintosh computers), and whether you want the scanning method to be Fast or Accurate, as Figure 1 shows. There's no option for Both. The software states that selecting Fast may skip some computers that are connecting wirelessly and selecting Accurate may take longer to resolve names. Although it's possible to simply skip this setup screen, for my test, I opted to scan SNMP devices and Macs, and I selected Accurate for the scan method.

After completing this screen, the software launched and immediately began scanning the subnet I was on. I didn't want this to happen at all, and this is my main complaint with the product. While I was obviously going to scan something, I wanted the ability to select what I wanted to scan first and not have the software start scanning for me.

After stopping this automatic scan, I found it easy to select the assets that I wanted to scan either by name or by IP address range. The software uses the Fluent UI pioneered by Microsoft with Office 2007, so I felt right at home. I provided administrator-level credentials as the software suggested and let it scan the handful of assets I had selected. After the software finished its scans, I began to explore the inventory details that it was able to capture.

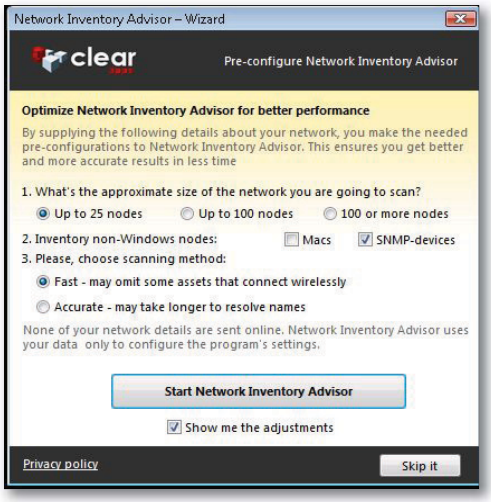

Figure 1: Network Inventory Advisor's pre-scan configuration settings

One thing that stood out for me was how comprehensive the scan was. On my Windows Vista test machine, the software detected basic details, such as OS version, and also details that I didn't expect to see, such as the date Windows was installed, product keys and serial numbers for Windows and Office, and visible and hidden Windows shares. I was also impressed to see details that other products don't usually bother to include, such as the configuration of PCI slots, the BIOS date and version, and the system's serial number. In fact, the software can be configured to scan for even more details, such as the Windows services that are installed.

Reporting on these collected inventory details is nice and easy to configure. Lots of predefined reports are available based on the hardware details (e.g., RAM, manufacturer) or based on the software details (e.g., OS, services installed, antivirus installed). If the included reports don't suit your needs, you can create your own.

I have another complaint about the product, aside from the automatic scan on startup. It correctly detected that Windows Defender was disabled, and it warned me about this situation. However, it didn't realize that this warning was unnecessary, even though it had also detected that I had Forefront Client Security (FCS) installed and that it was providing antivirus and anti-malware services.

Overall, despite my few complaints, I was impressed with Network Inventory Advisor. Pricing is available to match the size of your network, but even the site-wide pricing won't make you scream in terror. Take advantage of the 15-day trial and see if it works for you.

InstantDoc ID 129023

ClearApps Network Inventory Advisor

PROS: Easy to use; nice reporting; captures details other products don't

CONS: Automatic scan on first run should be a choice; antivirus detection isn't as smart as I hoped

RATING: ◆◆◆◆◇

PRICE: Under $1,000 for a site-wide license

RECOMMENDATION: Network Inventory Advisor fits the bill for a comprehensive product that won't drain your budget.

CONTACT: ClearApp • 916-509-7292 • www.clearapps.com

Michael Dragone | mike@mikerochip.com

REVIEW

Blackbird Management Suite

When it comes to managing your Active Directory (AD) environment, you need to make sure that you have a clear picture of where you've been and where you're going. **Blackbird Management Suite** helps you do that, letting you know exactly what's going on.

The installation takes place on a standalone server with a connection to either a separate Microsoft SQL Server computer or a local SQL Server Express database. A quick start guide walks you through configuring the services, connecting to AD, and setting up the database. The installation process is a breeze.

After you install the Blackbird server, the Blackbird console is installed. Then, you install any additional packages that you purchased, such as Auditor for AD, Privilege Manager, Recovery Manager, or Protector. It's important to remember that each domain controller (DC) needs to have a data handler installed. This is done in the Management Suite. This would be a simple process if the data handler installation routine were in the .msi format; a simple Group Policy object (GPO) linked to the DC would ensure that any new DC would always have Blackbird support. Unfortunately, the software is in the .exe format, so you have to remember to deploy the data handler to each DC that you add.

Blackbird Management Suite can perform a number of functions for AD, including recovery, auditing, protection from unwanted changes, and change management with workflow. I spent some time in each module and enjoyed the easy-to-understand interface as I worked to grasp the power of this suite of applications.

After I used the quick guide to set up the suite, it led me through setting up my first collector. A collector is simply a way to selectively back up the domain. While it may be desirable in a small company to back up the entire domain, an administrator of a large AD domain that spans multiple continents may want to compartmentalize the backups. Creating a collector is as simple as a few clicks. However, I think it could be even easier. The quick guide instructs you to "click the browse button" to select the OU or container that you want to back up. Instead, I found that you must manually enter the full path yourself (e.g., CN=Users,DC=itpro,DC=local). Regardless, to use a collector backup to restore a deleted object in AD, simply right-click an object in the domain, click Rollback, and select the backup that you want to use.

Of course, unexpected or unauthorized changes shouldn't occur in the first place, and that's where the Active Directory Rules feature comes in. With this feature, you can intervene when an object is created, modified, deleted, moved, or renamed. When one of these actions occurs, Blackbird can force the user to conform to a specific naming standard, prevent the action from occurring in the first place, submit the action for approval, send an email message, or run a script that you wrote. For example, perhaps you want to be alerted when a new user object is created. Or, maybe you want to approve all SMTP email address changes before they are implemented. I used Active Directory Rules to create quite a few rules, and I was impressed with its range and flexibility.

At the end of the day, you probably need a quick and dirty report on what's going on within your AD domain. Blackbird provides an audit trail so complete that a rogue administrator would be hard-pressed to execute his dastardly scheme. I counted no fewer than 43 built-in audit reports, not including the 18 preconfigured FISMA/HIPPA/PCI/SOX reports. If the included reports don't fit your particular need, you can easily create your own or modify an existing report.

In addition to the Blackbird console, the application is tightly integrated with Active Directory Users and Computers. By right-clicking an object, you can roll back a change, show a complete audit trail, or display account activity. Figure 1 shows the simple Blackbird interface and the tight integration with Active Directory Users and Computers. Selecting *Show account activity* generates a screen that displays a history of what the account has been doing—creating users, moving objects to other OUs, deleting computers, modifying object attributes, and so on.

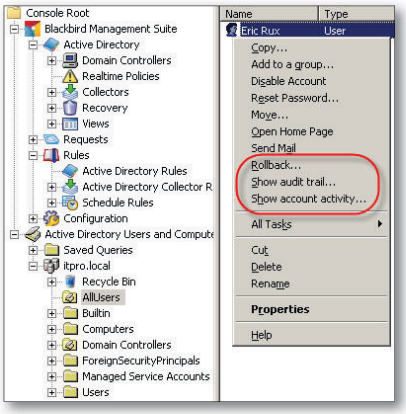

Figure 1: Blackbird's integration with Active Directory Users and Computers

There are a few features that could use some tweaking (e.g., adding a browse button to point to OUs instead of requiring the user to manually type the OU path), but I found the product fun to use. The Activity report should be built in to Windows Server by default! It makes finding out who did what as easy as a right- click. If you need some additional control over your Windows domain, add Blackbird to your list of products to evaluate.

InstantDoc ID 129013

Blackbird Management Suite

PROS: Easy to set up; easy to use; simple navigation with an intuitive interface

CONS: Some settings require manual entry; interface appears to be missing a "browse" button; modules can be confusing

RATING: ◆◆◆◆◆

PRICE: $14.40 per user for all four modules; individual prices and volume discounts available.

RECOMMENDATION: I highly recommend this easy-to-use product for all your AD management needs.

CONTACT: Blackbird Group • 866-224-8330 • www.Blackbird-group.com

Eric B. Rux | ebrux@whshelp.com

P2V Conversion Tools

Head-to-head comparisons of physical-to-virtual migration products from VMware, Novell, and Quest Software

by Michael Otey

Nowadays in many organizations, virtualization is the norm, and you need to have a good reason not to virtualize a given server. When it comes to virtualizing existing servers, you essentially have two choices: You can rebuild the server from scratch or you can use a Physical-to-Virtual (P2V) tool to automatically convert existing physical servers to virtual machines (VMs). Each method has its advantages.

Most administrators prefer to rebuild servers from scratch whenever they can. This tends to be more reliable and prevents the inevitable registry corruption that occurs in Windows systems. However, that method is definitely more time and labor intensive. In addition, it opens the door to the possibility of omitting required components. In contrast, P2V is faster, the conversions can be automated, and the end result contains all the applications and data from the source system. However, any problems that may have been present in the source system will also be transferred to the target system. For many IT professionals, the savings in time and effort make these P2V products well worth the investment and trade-off. In this review, I'll compare three of the leading P2V products: VMware vCenter Converter, Novell's PlateSpin Migrate, and Quest Software's Quest vConverter.

The P2V Process

The P2V process takes the physical system and converts it to a virtual machine specification and to one or more virtual hard disks. However, you can't just convert a physical hard disk to a virtual hard disk and expect the system to run. The P2V process must replace the hardware device drivers that are used in the physical machine with new device drivers that will work with the targeted virtualization platform. It's worth noting that this "hardware" change often results in the need to reactivate the converted OS.

In reviewing these tools, I considered the range of their support for the current virtualization platforms. I also assessed each product in terms of its viability in performing P2V conversions, in converting between different virtual platforms, in customizing the guest image that is being converted, and in automating the conversion process.

VMware vCenter Converter

VMware offers both a free standalone version of its vCenter Converter product and a version that is a plug-in to vCenter Server. The primary difference is that the standalone converter is used on a pay-per-support-incident basis, while support for the integrated vCenter Converter is included with the vSphere license. Additionally, the integrated vCenter Converter supports cold-cloning using a boot CD, and it can use a scheduler for centrally managing recurring conversions. In this review I looked at the standalone vCenter Converter. You can find a detailed comparison of the standalone and integrated vCenter Converter products at www.vmware.com/products/converter/get.html.

vCenter Converter can convert physical computers running 64-bit versions of Windows XP, Windows Server 2003, and Windows Server 2008, and 32-bit versions of Windows NT SP4 and later, Windows 2000, Windows XP, Windows Server 2003, Windows Server 2008, and Linux (RHEL, SUSE, and Ubuntu). It also supports conversions from most of the popular VM formats, including all VMware VMs, Microsoft Virtual PC, Microsoft Virtual Server 2005, Microsoft Hyper-V VMs (imported as a physical source), and Parallels Desktop. Additionally, vCenter Converter supports the following backup image formats: Symantec Backup Exec, Norton Ghost, Acronis True Image, and StorageCraft. However, one important limitation with VMware Converter is that it supports only VMware VMs as a destination. The VMs created by vCenter Converter can be used by VMware Workstation, VMware Player, VMware ESX/ESXi, VMware Server, and VMware Fusion.

Installation of vCenter Converter was fast and easy. During the installation, you can choose between installing Converter on the local system or in a client/server configuration. The local installation option runs all conversions on the local machine. The client/server option lets you run conversions from a remote system while you manage them from a local client system. I installed vCenter Converter by using the client/server configuration. This involved running setup on a central server to install the server components, and then running the installation on a network client to install the management console. I had to manually open port 443 to connect the server and the client. The standalone vCenter Converter supports both local and remote migrations if it's installed in a

client/server configuration. It also supports hot-cloning of multiple simultaneous conversions.

The built-in Conversion Wizard made it easy to create migration jobs. The wizard steps you through the process of specifying the source machine. Then it automatically deploys the conversion agent to the remote machine and prompts you for the name of the target ESX server. Unlike PlateSpin Migrate, which lets you choose different output targets, vCenter Converter is limited to outputting files onto a target ESX server. Next, the wizard lets you customize the VM that is created on the target. However, it doesn't let you change guest OS properties such as the OS name. After you complete the steps in the wizard, migration begins automatically. You can see vCenter Converter's UI in Figure 1.

Migrations can take some time, depending on the size of the system you are converting, the network speed, and the speed of the storage subsystem. In these tests on my 1GB network, a Windows system with a 300GB hard drive took about five hours to convert. At the end of the process, the Windows systems booted right up with no problems, and the Windows OS automatically recognized the new virtual device drivers. However, in keeping with the hardware changes, the system did require reactivation. I had no problems with any of the test migrations I performed using vCenter Converter.

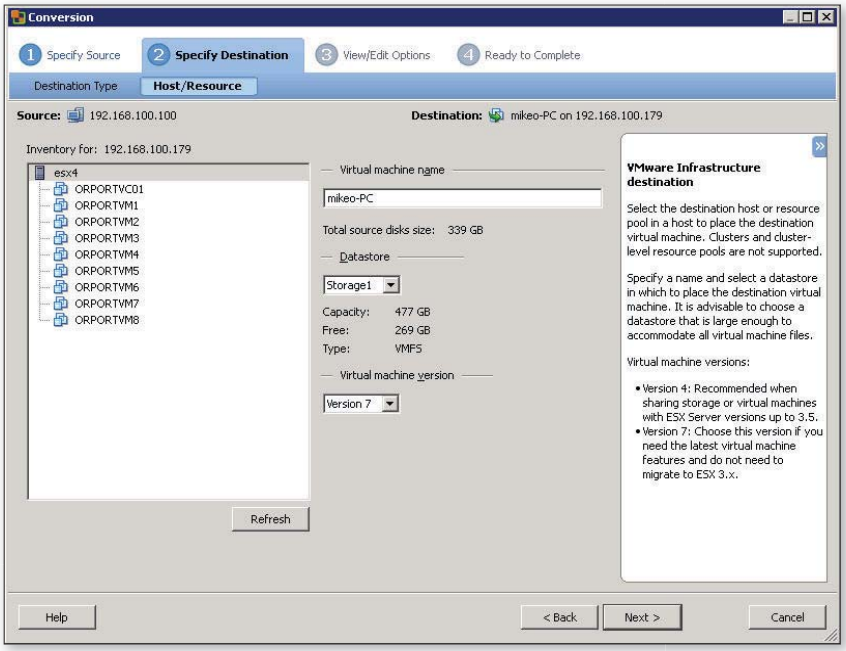

Figure 1: VMware vCenter Converter's UI

VMware vCenter Converter

PROS: Very easy to use

CONS: No support for Hyper-V or other non-VMware targets

RATING: ◆◆◆◇◇

PRICE: Standalone version is free; support is $90 per incident

RECOMMENDATION: Good for VMware-only conversion with a minimum of customization

CONTACT: VMware • 877-486-9273 • www.vmware.com

PlateSpin Migrate

Once a standalone company, PlateSpin was purchased by Novell in 2008. Novell now offers a number of PlateSpin products. PlateSpin Migrate is a more comprehensive product than vCenter Converter, but it's also more complex. While vCenter Converter is oriented toward converting physical and virtual systems into VMware VMs, PlateSpin Migrate is marketed as platform-agnostic—it converts physical systems to multiple virtualization formats (P2V). It also can convert VMs between multiple virtualization formats (V2V). Or it can go the other way and convert virtual machines to physical machines (V2P). PlateSpin Migrate is also able to capture and deploy server images. You can perform multiple migrations manually or you can schedule them. PlateSpin Migrate is also adept at keeping servers in sync for disaster recovery scenarios.

PlateSpin Migrate supports the following server operating systems: Windows Server 2008 (32- and 64-bit), Windows Server 2003 (32- and 64-bit), Windows 2000, Windows NT 4, SUSE Linux Enterprise Server (32- and 64-bit), Red Hat Linux (32- and 64-bit), and Sun Solaris. Additionally, you can use PlateSpin Migrate to convert the following desktop operating systems: Windows Vista (32- and 64-bit), Windows XP Professional, and Windows 2000. PlateSpin Migrate supports most of today's virtualization platforms as both source and target conversions. It also supports VMware ESX/ESXi, Microsoft Hyper-V, Citrix XenServer, SUSE Linux Enterprise with Xen, and Solaris Containers. Finally, PlateSpin Migrate supports the following image formats: Acronis TrueImage, PlateSpin Flexible Image Packages, Symantec Ghost, Symantec LiveState, Symantec Backup Exec, CommVault, and Tivoli Storage Manager.

Unlike vCenter Converter, PlateSpin Migrate requires a multi-server installation. On the backend, the product uses Internet Information Services (IIS) and a SQL Server 2005 database. It can use an existing SQL Server 2005 instance, or it can install an instance of SQL Server 2005 Express. As its name suggests, PlateSpin Migrate Server's task is to perform the migrations. You manage PlateSpin Migrate Server by using PlateSpin Migrate Client. I found the installation process for PlateSpin Migrate Server to be very difficult. First, the launcher always required that I expand its files before it presented the installation screen, thereby introducing an annoying delay into the setup experience. Next, the server itself had numerous dependencies. For example, it required 32-bit IIS, which wasn't mentioned in the requirements. Not only that, but the PlateSpin Knowledge Base (KB) instructions about enabling 32-bit IIS didn't work on my 64-bit Windows Server 2008 system. Therefore, I had to drop back and perform the installation on an actual 32-bit system. Not a very desirable option.

It took the better part of a day running and rerunning the setup program to load and install all the prerequisites such as IIS and the .NET Framework. The installer prompted me only when something that was needed wasn't present. Moreover, the product required multiple activations. This was most certainly one of the worst install experiences I have ever had. For a product with this many prerequisites, I suggest they look at the way Microsoft handles this situation for SQL Server 2008, where a screen in the installation process checks for all dependencies and then reports on whether the system meets these dependencies (and if not, gives you a heads-up on what you need to do to make the system ready).

That said, PlateSpin Migrate did not require agents to be installed on the source or target servers. Additionally, the user's guide was complete and helpful and was definitely a requirement for working with the product. To start, I had to register servers by supplying the DNS name or the IP address and authentication information. The servers were then registered in the client, as you can see in Figure 2.

PlateSpin Migrate functioned well for Microsoft Virtual Server 2005 and for VMware's ESX Server as both source and target. To get started, I needed to use the Discover Details options to catalog the systems I wanted to migrate. While it required no agent, I found the Discover Details process unreliable. It often reported network or authentication errors where there were no real problems. After discovery, you use the Conversion Wizard to define the migration tasks that you want to perform. The Conversion Wizard provides a great deal of control over the entire migration process. It walks you through steps such as providing credentials to access the source and target servers, selecting file or block transfers, performing live or offline migrations, providing a new host name, changing the networking configuration, and accessing the target's storage and services. For even more explicit control, an advanced peer-to-peer conversion option enables access to all migration options. I had no trouble running P2V migrations with VMware as the target, and the PlateSpin conversion ran smoothly, performing fully live migrations with no interruption in the source system's services. PlateSpin Migrate converted a 20GB server to a VM in about an hour and half. A new VM was created and started on the target platform without any problems.

While P2V worked great with VMware ESX Server as the target, I had trouble running V2V because PlateSpin Client wouldn't work with Hyper-V virtual machines in its Copy or Move Workloads dialog boxes. I also found PlateSpin Migrate to have limited Hyper-V support as a target. Hyper-V VMs weren't shown in the PlateSpin management console per virtualization host the way they were for Virtual Server 2005 or ESX Server. Additionally, while Hyper-V was supported as a target migration platform by PlateSpin Migrate, I had to perform numerous manual steps, such as creating the VM and supplying an ISO image for booting, before performing migrations.

PlateSpin was clearly the most feature-rich product in this review. However, it was also by far the most difficult to set up, and its implementation was stronger for VMware than for Hyper-V. I wouldn't recommend this product for Hyper-V migrations, but it performed well in P2V migrations to VMware's ESX Server.

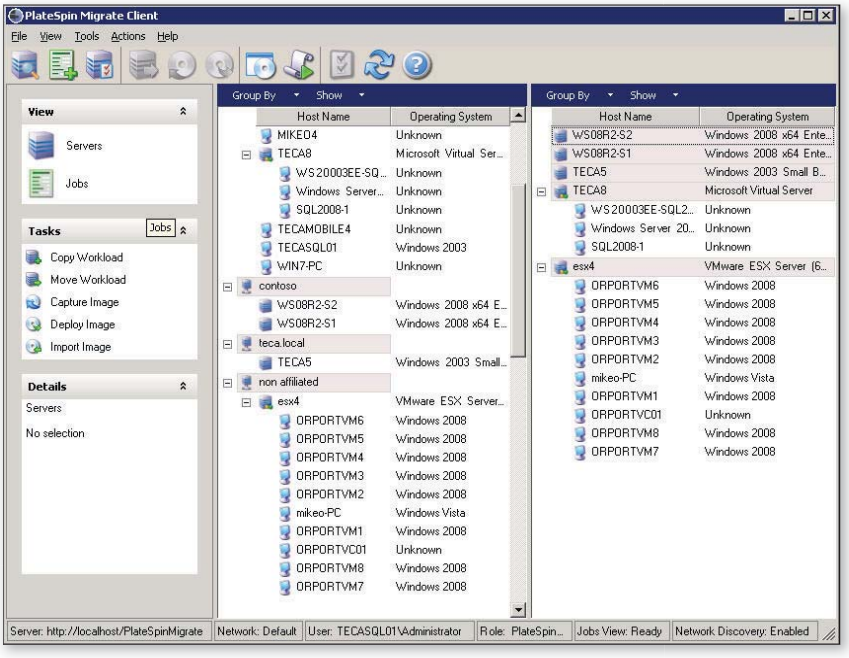

Figure 2: PlateSpin Migrate Client's UI

Novell PlateSpin Migrate

PROS: Converts to and from all popular virtualization formats

CONS: Poor support for Hyper-V

RATING: ◆◆◆◇◇

PRICE: $295 per workload

RECOMMENDATION: Good for VMware conversions that require significant customization. Not recommended for Hyper-V.

CONTACT: Novell • 800-529-3400 • www.novell.com

Quest vConverter

Quest Software's Quest vConverter can work with Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows XP Professional, and Windows Vista Ultimate as source systems. It can target both Microsoft Hyper-V and VMware's ESX Server hypervisors.

Quest vConverter is a simpler product to install than either vCenter Converter or PlateSpin Migrate. There is no multi-server option. Instead you install vConverter on the system on which you intend to run the conversions. vConverter can run on Windows XP Professional, Windows Vista Ultimate, Windows Server 2003, and Windows Server 2008 systems. The installation is a simple process, and there are no requirements for other auxiliary server products. vConverter uses Distributed Component Object Model (DCOM), and it requires you to open port 135 if there is a firewall between the vConverter system and the source systems.

vConverter tries to perform almost all of its management functions by using a single window, and this approach was not

completely successful for me. You can see the vConverter management console in Figure 3.

I found vConverter's UI to be clumsy and somewhat unprofessional-looking compared to vCenter Converter's and PlateSpin Migrate's UIs. The icons look dated and the screen's panes act like they should resize, but they don't. The Task List, in particular, was way too small. In spite of its name, the included Network Browser was not able to detect networked virtualization hosts or client systems when I tried it, but it did provide an option for importing system names from a .csv file. To get started, I manually entered the names of my Hyper-V and ESX servers.

You can perform P2V, V2V, and V2P migrations by running the Conversion Wizard. The Conversion Wizard was easy to use, but it didn't offer anywhere near the number of options and customizations provided by PlateSpin Migrate. The wizard first prompts you to choose between a P2V or V2P migration. The P2V option can be used for both P2V and V2V migrations. Next, you select the source and target servers and provide the required authentication information. You also select the server's target folder, and you can opt to change the VM's name. You can't set any of the detailed system properties, such as the system's name or storage configuration, as you can in PlateSpin Migrate. However, you can adjust the VM's properties, including the networking configuration. A Live Log tab lets you track the progress of your migration tasks.

I successfully used vConverter for both P2V and V2V migrations, with both ESX Server and Hyper-V as the target. All the conversions that used ESX Server as the target worked fine. However, some of the Hyper-V conversions of Windows Vista systems triggered blue-screen errors when the VM was started. All my conversions of Windows Server 2008 and Windows Server 2003 systems worked fine.

vConverter supports more than one-time migrations. Like PlateSpin Migrate, it also offers a VM synchronization capability, which is called Continuous Protection mode. In Continuous Protection mode, the target VM can synchronize changes with the physical system. Continuous Protection mode sends the changes only from the physical system to the target VM.

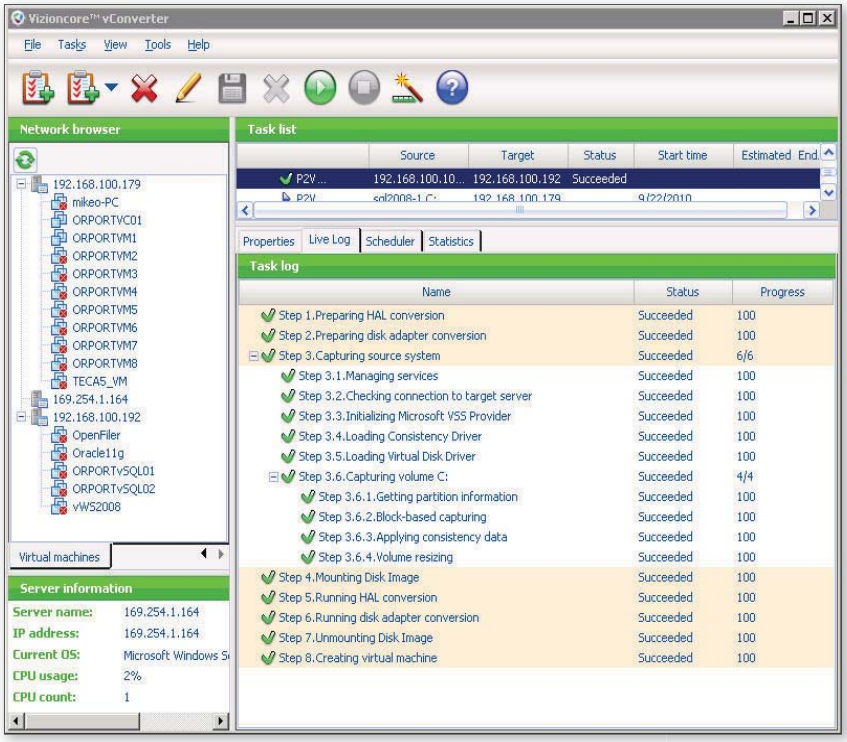

Figure 3: Quest vConverter management console

Quest vConverter

PROS: Very easy to install; performs P2V and V2P; supports both ESX Server and Hyper-V

CONS: VMware-only targets

RATING: ◆◆◆◆◇

PRICE: $299 per server

RECOMMENDATION: Good all-around choice for mixed P2V and V2V VMware and Hyper-V migrations.

CONTACT: Quest Software • 866-260-2483 • vizioncore.com

Something Old and Something New

These P2V tools take something old—your physical systems—and turn them into something new: virtual machines. These tools are useful for helping you consolidate your server infrastructure. Be prepared for the migration process to take several hours, depending on your infrastructure and the size of the servers that are being converted. Even so, the process is faster than building a new server from the ground up and then restoring the local data. These tools let you automate the conversions—you can set them up and walk away. Although some experts suggest that these types of products are suitable for backup and disaster recovery, my experiences indicate that the technology is good for mass migrations but not reliable enough to bet your business on it in disaster recovery scenarios.

Of the three products reviewed here, Quest's vConverter was my clear pick for editor's choice. It was the only product that worked well for both VMware ESX Server and Hyper-V. vCenter Converter was the easiest to use, but its VMware-only orientation made it far more limited than the other products. PlateSpin Migrate had the most powerful migration customization capabilities and control, but its arduous setup process, unreliable interface, and poor Hyper-V support made it less useful than Quest's simpler vConverter, which offers better support for Hyper-V. If you have a number of migrations to perform, these products can certainly help. However, this is definitely a technology area that needs more time to mature.

InstantDoc ID 129102

Michael Otey
(motey@windowsitpro.com) is senior technical director for *Windows IT Pro* and *SQL Server Magazine* and author of *Microsoft SQL Server 2008 New Features* (Osborne/McGraw-Hill).

KVM Over IP

Improve your—and your data center's—efficiency

by Jason Bovberg

Particularly if you administer a sprawling IT environment, you face the challenge of overcoming geographic barriers in your day-to-day network management; you need to quickly react to problems on far-reaching systems. Or if—in the clutches of our weak economy—you're all alone, performing solo IT administration at a financially strapped small-to-midsized business (SMB) or branch office, you need to increase productivity despite your lack of resources. Either way, the KVM over IP switch is probably one of the most essential components in your network infrastructure. This buyer's guide provides a sampling of the market's current offerings.

KVM Benefits

It's clear how the functionality of such a device can improve your day-to-day efficiency: A KVM over IP switch gives you in-band or out-of-band access to system keyboard, video, and mouse (KVM) functions, from any location at any time.

A KVM over IP switch lets you maintain and manage geographically diverse devices, better manage systems to deliver key business services, and reduce total cost of ownership (TCO). KVM over IP switches give you BIOS-level control of connected servers and other network devices straight from any location: From a central interface, you can securely manage your entire IT infrastructure—including branches and remote data centers—as if you were administering them locally. A good KVM switch gives you complete access to authentication, event alerts, and user log files. Some KVM solutions even let you manage all your servers and devices when the network has failed and remote-access software isn't functioning.

Factors to Consider

KVM over IP switches can differ substantially in their breadth of functionality. To avoid wasting valuable resources or even compromising your business's security, you need to consider carefully the options you need for your unique environment. For example, the solution you choose needs to be able to support every OS platform and network device in your environment. The solutions in Table 1 support a broad range of platforms. You might not have some of these platforms in your local environment, but don't forget that your network probably knows no boundaries; you must also consider remote users' laptops and mobile devices.

How many ports do you want the switch to have? As your company inevitably grows following this downturn, you'll need it to handle more than it needs to handle now. Switches differ wildly in the number of computers that can connect to it, and in enterprise scenarios, you can daisy-chain switches to cover greater numbers of connections. How does the switch handle video? What's the maximum video resolution? For bandwidth conservation, check to see what type of video compression the switch offers. Another feature you might find useful is sound capability. What about the switch's form factor (is it rack-mountable?), the type of cables you'll need for server connections, the maximum number of simultaneous sessions, and the maximum distance the switch allows between the switch and servers? And what kind of failover functionality does the switch provide? Reliable access to critical resources is a key feature of a KVM over IP platform.

Some switches offer proprietary viewer software for communicating with the KVM switch, whereas others rely on a web browser to perform the same function. If you prefer limited user access to the switch, client software might be best. But see if you can get a handle on usability and performance; entry-level products might offer relatively weak security and reliability. If you need to give administrators access regardless of location, go with the browser-based interface.

Speaking of security, a major byproduct of the KVM over IP switch's inherent centralization is tighter control of your widespread resources, but the various solutions available today take differing approaches to security. Determine whether the switch takes advantage of your existing authentication technologies or uses its own methods. Does the switch encrypt all signals between itself and managed devices? A great deterrent to intrusion is an encrypted administrative GUI.

Choose Wisely

Table 1 shows a listing of the vendors who chose to participate in this year's roundup of KVM over IP switches. You might consider KVM technology basic or elemental, but it's one area where you don't want to choose unwisely.

InstantDoc ID 129145

Jason Bovberg
(jbovberg@windowsitpro.com) is a senior editor for Windows IT Pro, SQL Server Magazine, and System iNEWS, specializing in networking, hardware, storage/backup, and mobile and wireless. He has 20 years of experience as a writer and editor in magazine, book, and special-interest publishing.

Company	Product Name	Price	Form Factor	Dimensions
Adder 978-499-2105 888-932-3337 www.adder.com	AdderView CATxIP 5000	$1,750	Desktop/1U rackmount	7.92" x 1.76" x 4.8"
	AdderView CATxIP 1000	$1,195 (8 port), $1,395 (16 port)	Desktop/1U rackmount	7.92" x 1.76" x 4.8"
	AdderView CATxIP	$2,395 (16 port), $2,695 (24 port)	Desktop/1U rackmount	19" x 1.8" x 8.5"
	AdderLink IPEPS	$449	Desktop/2U rackmount	4.72" x 1.65" x 2.95"
	AdderLink IPEPS DA	$699	Desktop/2U rackmount	4.72" x 1.65" x 2.95"
	AdderLink IP	$995	Desktop/1U rackmount	7.92" x 1.76" x 4.8"
	AdderLink IP Gold	$1,495	Desktop/1U rackmount	7.92" x 1.76" x 4.8"
Belkin 310-751-5100 www.belkin.com/kvm/sms	OmniView 5xxxK Series	From $399	1U rackmount	17" x 10.6" x 1.7"
Lantronix 949-453-3990 www.lantronix.com	SpiderDuo	$385	Not rackmounted	5.2" x 2.3" x 1.4"
StarTech.com 519-455-9675 800-265-1844 www.startech.com	SV1654DX4I 4 User 16 Port Cat5 Matrix IP KVM Switch	$3,814	1U Rackmount steel chassis	17.3" x 12.2" x 1.7"
	CABCONS1716I 1U 17" Rack Mount LCD Console with Integrated 16 Port IP KVM Switch	$2,520	1U Rackmount steel chassis	17.3" x 12.2" x 1.7"
	SV3254DX4I 4 User 32 Port Cat5 Matrix IP KVM Switch	$4,200	1U Rackmount steel chassis	17.3" x 12.2" x 1.7"
	SV841HDIE 8 Port Rack Mount USB PS/2 Digital IP KVM Switch	$822	1U Rackmount steel chassis	8.8" x 15.9" x 1.7"
	SV3253DXI 32 Port Multi-User Cat5 Matrix IP KVM Switch	$2,780	1U Rackmount steel chassis	17.3" x 12.2" x 1.7"
	SV441DUSBI 4 Port USB VGA IP KVM Switch with Virtual Media	$729	Steel chassis/optional 1U rackmount kit	8.8" x 7.3" x 1.7"
	SV1653DXI 16 Port Multi-User Cat5 Matrix IP KVM Switch	$2,050	1U Rackmount steel chassis	17.3" x 12.2" x 1.7"
	SV1641HDIE 16 Port Rack Mount USB PS/2 Digital IP KVM Switch	$1,088	1U Rackmount steel chassis	8.8" x 15.9" x 1.7"
	SV1115IPEXT 1 Port USB PS/2 Server Remote Control IP KVM w/Virtual Media & Serial	$730	Steel case/optional 1U rackmount kit	0.9" x 6.6" x 7.4"
	SV1107IPEXT 1 Port Server Remote Control IP KVM w/Virtual Media	$340	Small steel case/optional 1U rackmount kit	1" x 3.6" x 6.6"
Tripp Lite 773-869-1234 www.tripplite.com	B072-016-1-IP	$2,499	1U rackmount	1.72" x 17" x 9"
	B070-016-19-IP	$4,000	1U rackmount	1.72" x 17" x 27"
	NetDirector Cat5 IP KVM Switch series	$6,000 - $10,000	1U rackmount	1.625" x 17" x 16"
	B051-000	$950	Desktop/rackmount	1" x 7.9" x 3"
	B020-U08-19-IP	$3,999	1U rackmount	1.625" x 17" x 27"
	B020-016-17-IP	$6,479	1U rackmount	1.625" x 17" x 26.5"

	Platforms Supported	Maximum Video Resolution	Sound Capability	Maximum Distance Between Master Switch and CPU	Maximum Allowable Port Limit	Ports Available per Switch	External Adapters Required to Accomplish Multiplatform Support?
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200 at 60Hz	No	33'	256	16	Yes
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200 at 60Hz	Yes	30'	256	8 or 16	Yes
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200	Yes	150'	256	16 or 24	Yes
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200	No	15'	8	1	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200	No	15'	8	1	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200	No	15'	256	1	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200	Yes	10'	256	1	No
	Windows, DOS, UNIX/Linux, SunOS, USB	Remote	No	150'	200	Up to 32	No
	Windows, UNIX/Linux, MacOS, USB	1600 x 1200 at 60Hz	Yes	59"	8	Up to 8	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200 at 60Hz	No	164'	4	16	Yes
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX	1600 x 1200 at 85Hz (web); 1920 x 1440 (local)	No	50'	1	16	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200 at 60Hz	No	164'	4	32	Yes
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX	1600 x 1200 at 85Hz (web); 1920 x 1440 (local)	No	50'	1	8	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200 at 60Hz	No	164'	3	32	Yes
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200 at 85Hz (web); 1920 x 1440 (local)	No	15'	1	4	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200 at 60Hz	No	164'	3	16	Yes
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX	1600 x 1200 at 85Hz (web); 1920 x 1440 (local)	No	50'	1	16	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200 at 85Hz	No	50'	1	1	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1280 x 1024	No	15'	1	1	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200	No	100'	256	16	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1280 x 1024 (local) / 1600 x 1200 (remote)	No	100'	256	16	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1600 x 1200	No	164'	256 or 512	16 or 32	No
	Windows, DOS, UNIX/Linux, SunOS, IBM/AIX, HP-UX, USB	1600 x 1200	No	19'-25'	Dependent on KVM connected	1	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1280 x 1024 (local) / 1600 x 1200 (remote)	No	15'	504	8	No
	Windows, DOS, UNIX/Linux, SunOS, MacOS, IBM/AIX, HP-UX, USB	1280 x 1024 (local) / 1600 x 1200 (remote)	No	19'-25'	256	16	No

INSIGHTS FROM THE INDUSTRY

Windows Phone 7: Is Microsoft Presenting the Right Message?

For IT pros and developers looking to get the goods on what's going on with Microsoft's new Windows Phone 7 platform, the keynote address (this past November) for Exchange Connections (and Windows Connections, and DevConnections, etc.) was the place to be. Microsoft's Joe Belfiore, corporate vice president for Windows Phone Program Management, delivered his message to a full house of eager listeners, and Belfiore didn't disappoint. With little preamble, he hooked his Windows Phone 7 device to a PC so he could broadcast his smartphone screen to the hall's big display screens, and then he proceeded to demo just about every feature of the phone that you could probably think of.

In fact, Belfiore's talk went over time, and I heard more than one person afterward say it was too long—but I didn't notice anyone getting up to leave. Although there were no new announcements about the phone, if you haven't had a chance to see it demoed, this was a great exploration both of how the UI is set up as well as some of the integration of the platform's various features. He talked about the minimum hardware requirements that all handset makers will be held to, as well as what UI elements the OEMs and carriers are able to customize in their particular devices—and what they can't. I won't bore you with all those details, which I'm sure you can find elsewhere. What is to me most interesting in all of this is how what I saw demoed live onstage so vastly outshines what I'm seeing advertised and pushed by Microsoft.

Microsoft has made it clear that it's marketing to the consumer audience, and I don't have a problem with that. Hey, it worked for iPhone and Android; that's where the smartphone market is these days. (Which doesn't mean the company couldn't also talk up the enterprise security and control features, but that's probably another discussion.) But even with that consumer focus, I don't think they're getting the right message out there, or showing off the right features.

For example, one of the first things Belfiore showed was how you could take a Windows Phone 7 phone that was locked, and press the camera button to immediately wake it up to take a picture. Even if the phone requires a password to unlock, you can still get to the camera with this method, although not to other phone features or functions. If Microsoft put this feature in a TV commercial, every parent who ever missed a shot of a child because they couldn't get their camera app up quick enough would be in line to buy this thing.

The real problem, which could be a crucial problem for Windows Phone 7, is simply not being clear about what the benefits really are. Instead, we get commercials of people bumping into each other and dropping phones in toilets.

Belfiore also showed off the capabilities of Bing through mobile browser experience. I haven't been the least bit impressed with Bing on the desktop, but seeing what it does on the phone, I could become a convert. Combining voice search and location detection, along with a bit of logic to give the results most likely to be what you want, you can speak the name of a current movie and instantly get a listing of theaters with show times and links to buy tickets. Or if you're looking for a restaurant, you can get results in your vicinity, with hours, menus, contact info, and so forth—exactly the snapshot you're probably looking for when you do a search from your mobile device. Microsoft has been advertising this as a "decision engine" without doing the job of explaining what that means.

The real problem, which could be a crucial problem for Windows Phone 7, is simply not being clear about what the benefits really are. Instead, we get commercials of people bumping into each other and dropping phones in toilets. I mean, it's an enjoyable commercial, but it's a confusing message, particularly because it isn't clear that those phones in the commercial causing all the problems aren't Windows Phone 7 devices. A phone to save us from our phones? Yeah, but most people really like their phones and don't want or need to be saved from them. (This spoken from a recently self-avowed smartphone addict, remember—#smartphoneaddict.)

Well, I hope people take the time to see these devices for themselves when they begin selling in the US next week. I believe that no one device is going to be the best solution for everyone; we're just not going to see a dominant mobile OS in the foreseeable future. Which means Windows Phone 7 certainly has a chance to find an audience; it might stand a better chance if Microsoft could figure out what it was trying to sell about the thing.

—B.K. Winstead

What is Tethering, and Should You Care?

I've been hearing a lot about tethering lately. It's been around for years, but I've never given it much thought, probably because (a) I have a dumb phone that doesn't even have Bluetooth, and (b) because I use T-Mobile on my cell phone, which doesn't officially support tethering.

Anyway, I don't really blame myself. Trying to navigate the truths about tethering is like finding dry land in Waterworld. In any event, I set out to answer a few basic questions. Below are the answers.

What exactly is tethering and how does it work? In simple terms, tethering is connecting to the Internet on a computer(e.g., a laptop or netbook) through a mobile device. In other words, the phone is used as a modem, and can connect to the computer via USB or Bluetooth.

What does it cost (monthly) to tether? It varies by each carrier and each individual smartphone. But the prices are generally:

- AT&T: $20, after a $25 2GB data plan.
- Sprint: $15, after a data plan (prices vary).
- T-Mobile: T-Mobile doesn't officially support it.
- Verizon: $15, after a $30 5GB data plan.

These are US rates—international rates vary. Rates may also vary for individual devices.

Also, my understanding is that technically you do not need a tethering plan, but the carrier reserves the right to charge you extra for tethering if you don't have one.

What devices support tethering? The iPhone 4, newer Blackberry devices, and most Android smartphones support tethering. There are too many devices in this category to compile a comprehensive list, however, so I recommend running a Google search or talking to your provider about a specific device. Windows Phone 7 devices won't support tethering at launch, but should in the near future.

Is the performance when tethering acceptable, and for what tasks? I have asked several people that use tethering about this, and I've received mixed responses. The general consensus seems to be that tethering is great when you're in a situation where you can't get Wi-Fi, or don't feel comfortable connecting to an unsecured Wi-Fi connection and just need to perform a few simple tasks. However, in most cases, you'd probably just prefer to access the Internet on your phone, as you'll get better performance. (The extra connection to the laptop slows performance, and your laptop will generally eat up more bandwidth to run the same sites and applications.)

As far as exact numbers, I've seen a few sites suggest that 3G tethering can average about 1 Mbps download, assuming all is going well. But it certainly varies based on the network you have and the coverage you're getting.

Finally, there's a growing number of individuals who opt to pass on broadband Internet access at home, saving about $50/month, and just use tethering for home computing. Well, let's just say those individuals probably aren't playing World of Warcraft. It'd probably be fine for checking email, surfing the web, and watching a few YouTube videos (though those YouTube videos will eat up your data plan limit if you have one), but beyond that it's not going to cut it.

In summary, the pros of tethering are:

- Convenient Internet access on your computer when you don't want to access a public network or one isn't available
- Can use as an alternative to home broadband Internet in small doses

The cons are:

- Cost—tethering will net you an extra $15-20 per month on top of the smartphone data plan
- Connection speeds are, at best, a little slower than a smartphone's performance

> Tethering can be a good alternative to having your employees using public Wi-Fi when they're on the go.

For businesses: If you have a lot of mobile workers and can get a decent corporate rate, tethering can be a good alternative to having your employees using public Wi-Fi when they're on the go. But, don't expect it to change the face of mobile Internet access anytime soon.

—Brian Reinholz

DISCOVER WINDOWS IT PRO VIP

Windows IT Pro VIP is the perfect tool for the IT pro who knows that 15 minutes searching the Web is costing more than just time.

WINDOWS IT PRO is:

1. **Educational**—FREE eLearning courses and eBooks to keep your skills sharp
2. **Deep**—over 41,000 articles on DVD and online, some exclusively for VIP members
3. **Broad**—all the articles, solutions, and FAQs ever published in:
   Windows IT Pro
   SQL Server Magazine
   SharePointPro Connections
   DevProConnections
4. **Reliable**—every solution has been road-tested by our experts
5. **Impartial**—with technical editors who are shaping the industry
6. **Economical**—more than $1,000 of resources for less than $17* a month

Upgrade to VIP at windowsitpro.com/go/vip

* Rates vary outside the U.S.

Making "Remote" as Close as a Click Away

TeamViewer, maker of the remote-desktop sharing and file-transfer solution of the same name, offers new features for online support on Windows computers. Every user on a terminal server now receives a TeamViewer ID, allowing supporters to connect to different terminal server users and access their accounts remotely. Simultaneous connections to more than one user are now possible, and support can be given from the terminal server to any external computer.

TeamViewer also offers drag and drop file transfer. During a remote support session, files can be dragged to the partner screen for transfer and will appear on the desktop of the partner computer. TeamViewer also automatically detects additional monitor screens, in the case of multi-monitor users.

To use TeamViewer, you download and install it on your computer and fire it up. Your partner or user being supported at the remote computer then downloads TeamViewer QuickSupport and starts it. QuickSupport doesn't have to be installed and can be executed without Windows admin rights. You then ask your partner for the TeamViewer ID and enter it in the ID field, click "Connect to partner," enter the partner's password, and you are connected.

TeamViewer is secured by RSA Public/Private Key Exchange andAES-256 session encoding. With every start of TeamViewer, a new dynamic session password is generated, which prevents any permanent access.

TeamViewer actually has a "demo hotline" where you can call and see it in action with a connection established to your computer. Or you can go old-school and check out their website.

—Caroline Marwitz

Lync 2010 Looks to Improve the UC Conversation

At Exchange Connections in Las Vegas this past November, in addition to lots of information about Windows Phone 7, Microsoft was really getting the word out about the newly renamed Office Communications Server, now known as Microsoft Lync. Shaun Pierce, the general manager of the Lync Server Division, delivered a keynote about Lync, which included the expected live demo of the collaboration possibilities of the new version. In addition to the new end-user enhancements, Lync should be easier to implement than OCS—which is good news for IT pros.

One of the first new things most people will notice with Lync is that contact photos appear in your Lync client, which is essentially an updated version of Office Communicator. But it's a pretty extensive update, which also includes the ability to show your contacts in a view similar to Facebook or Twitter with recent activity displayed. However, I spoke to Microsoft's Harold Wong in the Expo Hall, who assured me that the social media aspects can easily be turned off if you're someone who already feels overwhelmed by that sort of information. Other new features in the UI include the ability to use a dial pad to dial (yes, apparently some people prefer this method over dial by name) and visual access to voicemail.

Pierce's keynote also highlighted many of the architectural changes in Lync that lead to things like better call routing and—hopefully—easier deployment. The big change here is a reduction in the number of mediation servers needed because a single mediation server can now talk to multiple PSTN networks. The complexity of OCS 2007 was undoubtedly a deployment blocker for many organizations, and even for those that did implement the product, it might have prevented them from getting full benefits from it. We'll see if Microsoft got that message and has truly managed to improve things with Lync.

Another major change, and one of the features Wong mentioned as most significant for IT pros, is that Lync is now supported for virtualization. Wong also listed better bandwidth control through policies, wide device support, and improved SIP trunking support as things IT pros could look forward to with the Lync 2010 release. He demoed the ease of switching devices, such as microphones or video, on the fly during a conference without the need to sign off and back on to have the new device recognized. Indeed, there's a lot to be excited about, and talking with Harold Wong about UC is likely to inspire anyone. At the time of this writing, Lync was expected to reach general availability on December 1.

—B.K. Winstead

Windows IT Pro Twitter Feeds to Follow

@windowsitpro - Get the latest article updates across the Windows IT Pro website.
@sqlservermag - Updates and giveaways from our sister publication, SQL Server Magazine.
@savvyasst - Related events, resources, and giveaways in the Windows IT space.
@michelecrockett - Updates from Michele Crockett, Editorial and Custom Strategy Director
@witproAmy - Updates from Amy Eisenberg, Executive Editor
@wincaroline - Updates from Caroline Marwitz, editor specializing in Active Directory and SharePoint
@breinholz - Updates from Brian Reinholz, editor specializing in training, certification, and mobility
@zacwiggy - Updates from Zac Wiggy, editor specializing in systems management, Windows OS, and virtualization
@bkwins - Updates from B. K. Winstead, editor specializing in Exchange, Outlook, and mobility
@thurrott - News Editor Paul Thurrott's Twitter feed

SharePointPro CONNECTIONS

Join the SharePoint Expert Community

SharePointPro Connections

provides real-world advice from professionals and peers who share their experience administering and developing in SharePoint.

SharePointPro Connections is the independent voice on SharePoint technology. Expert authors provide our community members with field-tested information to enable content and image management, collaboration, and workflow solutions tailored to business needs.

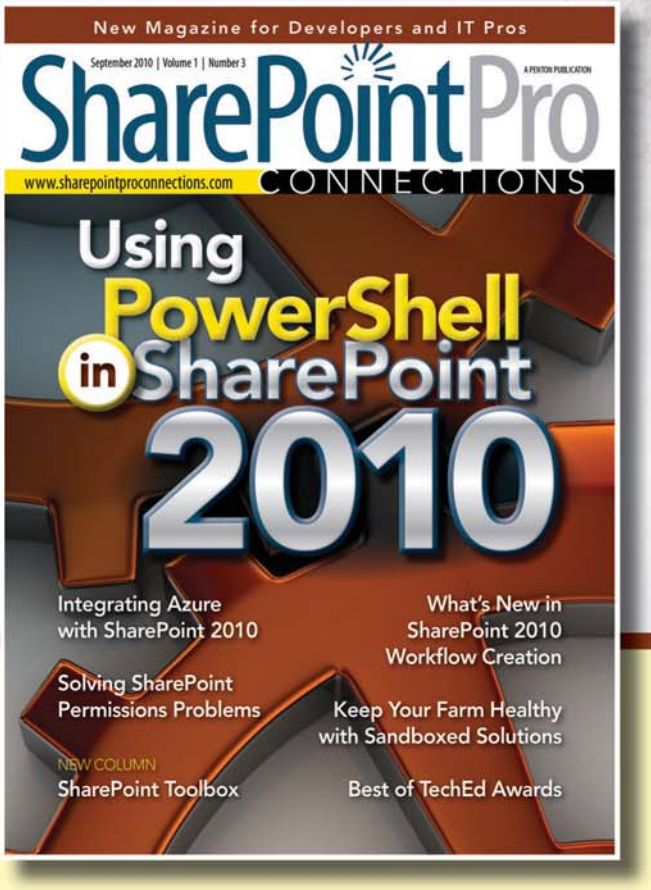

Upcoming topics include:

- Developing Line of Business Connectivity with SharePoint BCS (Business Connectivity Services)
- Planning for SharePoint in the Cloud and on Mobile Devices
- Deep Collaboration with SharePoint Social Media
- Developing and Administering SharePoint Business Intelligence (BI) Solutions

Subscribe FREE to the only magazine dedicated to all things SharePoint.

sharepointproconnections.com/go/SubscribeNow

AD INDEX

For detailed information about products in this issue of *Windows IT Pro*, visit the web sites listed below.

VENDOR DIRECTORY

The following vendors or their products are mentioned in this issue of *Windows IT Pro* on the pages listed below.

DIRECTORY OF SERVICES | WINDOWS IT PRO NETWORK

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowsitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
www.windowsitpro.com/go/forums

News

Check out the current news and information about Microsoft Windows technologies.
www.windowsitpro.com/go/news

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- *asp.netNOW*
- *DevProConnections UPDATE*
- *Exchange & Outlook UPDATE*
- *Security UPDATE*
- *SharepointPro Connections UPDATE*
- *SQL Server Magazine UPDATE*
- *Windows IT Pro UPDATE*
- *Windows Tips & Tricks UPDATE*
- *WinInfo Daily UPDATE*

www.windowsitpro.com/email

RELATED PRODUCTS

Custom Reprint Services

Order reprints of *Windows IT Pro* articles. Diane Madzelonka at Diane.madzelonka@penton.com.

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the Web. Includes FREE access to eBooks and archived eLearning events, plus a subscription to either Windows IT Pro or SQL Server Magazine.
www.windowsitpro.com/go/vipsub

SQL SERVER MAGAZINE

Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

ASSOCIATED WEBSITES

DevProConnections

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePointPro Connections

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and Web seminars mentored by a community of peers and professionals.
www.sharepointproconnections.com

NEW WAYS TO REACH *WINDOWS IT PRO* EDITORS:

LinkedIn: To check out the *Windows IT Pro* group on LinkedIn, sign in on the LinkedIn homepage (www.linkedin.com), select the Search Groups option from the pull-down menu, and use "Windows IT Pro" as your search term.

Facebook: We've created a page on Facebook for *Windows IT Pro*, which you can access at: http://tinyurl.com/d5bquf. Visit our Facebook page to read the latest reader comments, see links to our latest web content, browse our classic cover gallery, and participate in our Facebook discussion board.

Twitter: Visit the *Windows IT Pro* Twitter page at www.twitter.com/windowsitpro.

Windows IT Pro

by Jason Bovberg

Get Back to Work!

PRODUCT OF THE MONTH

Here's an extraordinarily useful gift for the office worker. The Blabber Meter is a desk clock that tallies the amount of money it costs your company for you to endure long meetings, conversations, and phone calls. "The inspiration in developing the Blabber Meter," writes product developer Brad Johnson, "came from working at a company that regularly held four-hour weekly staff meetings. At the end of each meeting, I walked away wondering what did that accomplish?" Or perhaps you—like us—are wary of actually visualizing the real cost of blabber. Those dollars can really pile up. The Blabber Meter takes two AAA batteries (included), costs $24.99, and is available at www.blabbermeter.com.

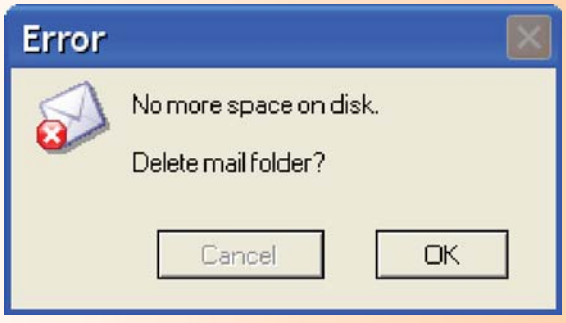

Figure 1: Sure, why not?

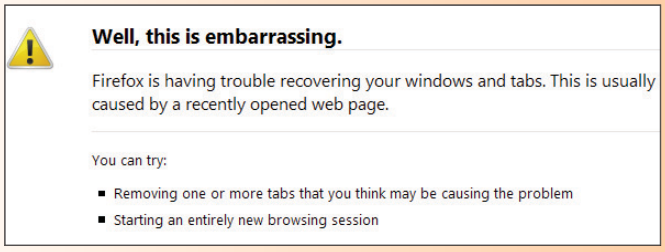

Figure 2: Red-faced fox

USER MOMENT OF THE MONTH

One of the funniest moments I ever experienced was back in the days of the diskette. I was walking a user through a remote install on his system. The user had received a series of diskettes containing the software, and he was halfway through the installation process when he said, "Hmm, I can't get disk 2 out of the drive." I asked him if the installation process had instructed him to remove the disk, and he said, "Yes." I asked him if he'd firmly pressed the Eject button. "Yes." I could hear him jabbing the button over the phone. I was calmly walking through possible scenarios, but I noticed a slight edge of panic in his voice, as if the machine had consciously betrayed him. I finally had to send the local support guy out to his desk an hour later. He called me a few seconds after arriving at the desk to say the user had ejected the disk and forgotten. It was resting in plain sight on top of the tower.

—Briane

January 2011 issue no. 197, *Windows IT Pro* (ISSN 1552-3136) is published monthly. *Windows IT Pro,* 221 E. 29th St., Loveland, CO 80538, (800) 793-5697 or (970) 663-4700. Sales and Marketing Offices: 221 E. 29th St., Loveland, CO 80538. Advertising rates furnished upon request. Periodicals Class postage paid at Loveland, Colorado, and additional mailing offices. **POSTMASTER: Send address changes to *Windows IT Pro,* 221 E. 29th St., Loveland, CO 80538. SUBSCRIBERS:** Send all inquiries, payments, and address changes to *Windows IT Pro,* Circulation Department, 221 E. 29th St., Loveland, CO 80538. Printed in the USA.

YOUR

REMOTE SUPPORT SILVER BULLET.

You've heard the call for help and are ready to ride – all you need is the best tool to let your skills shine. With the simple yet powerful technology of GoToAssist® Express™, you'll connect with customers like never before.

Speed – Instantly support up to 8 clients at once.
Unlimited Use – Wrangle all the issues you want for one flat fee.
Performance – 100% reliability you can hang your hat on.
Unattended Support – Work while customers are away.

GoToAssist® EXPRESS™

"One support slinger against a world of technical problems? I like my odds."

by CITRIX®

Try It FREE for 30 Days
www.gotoassist.com/ad

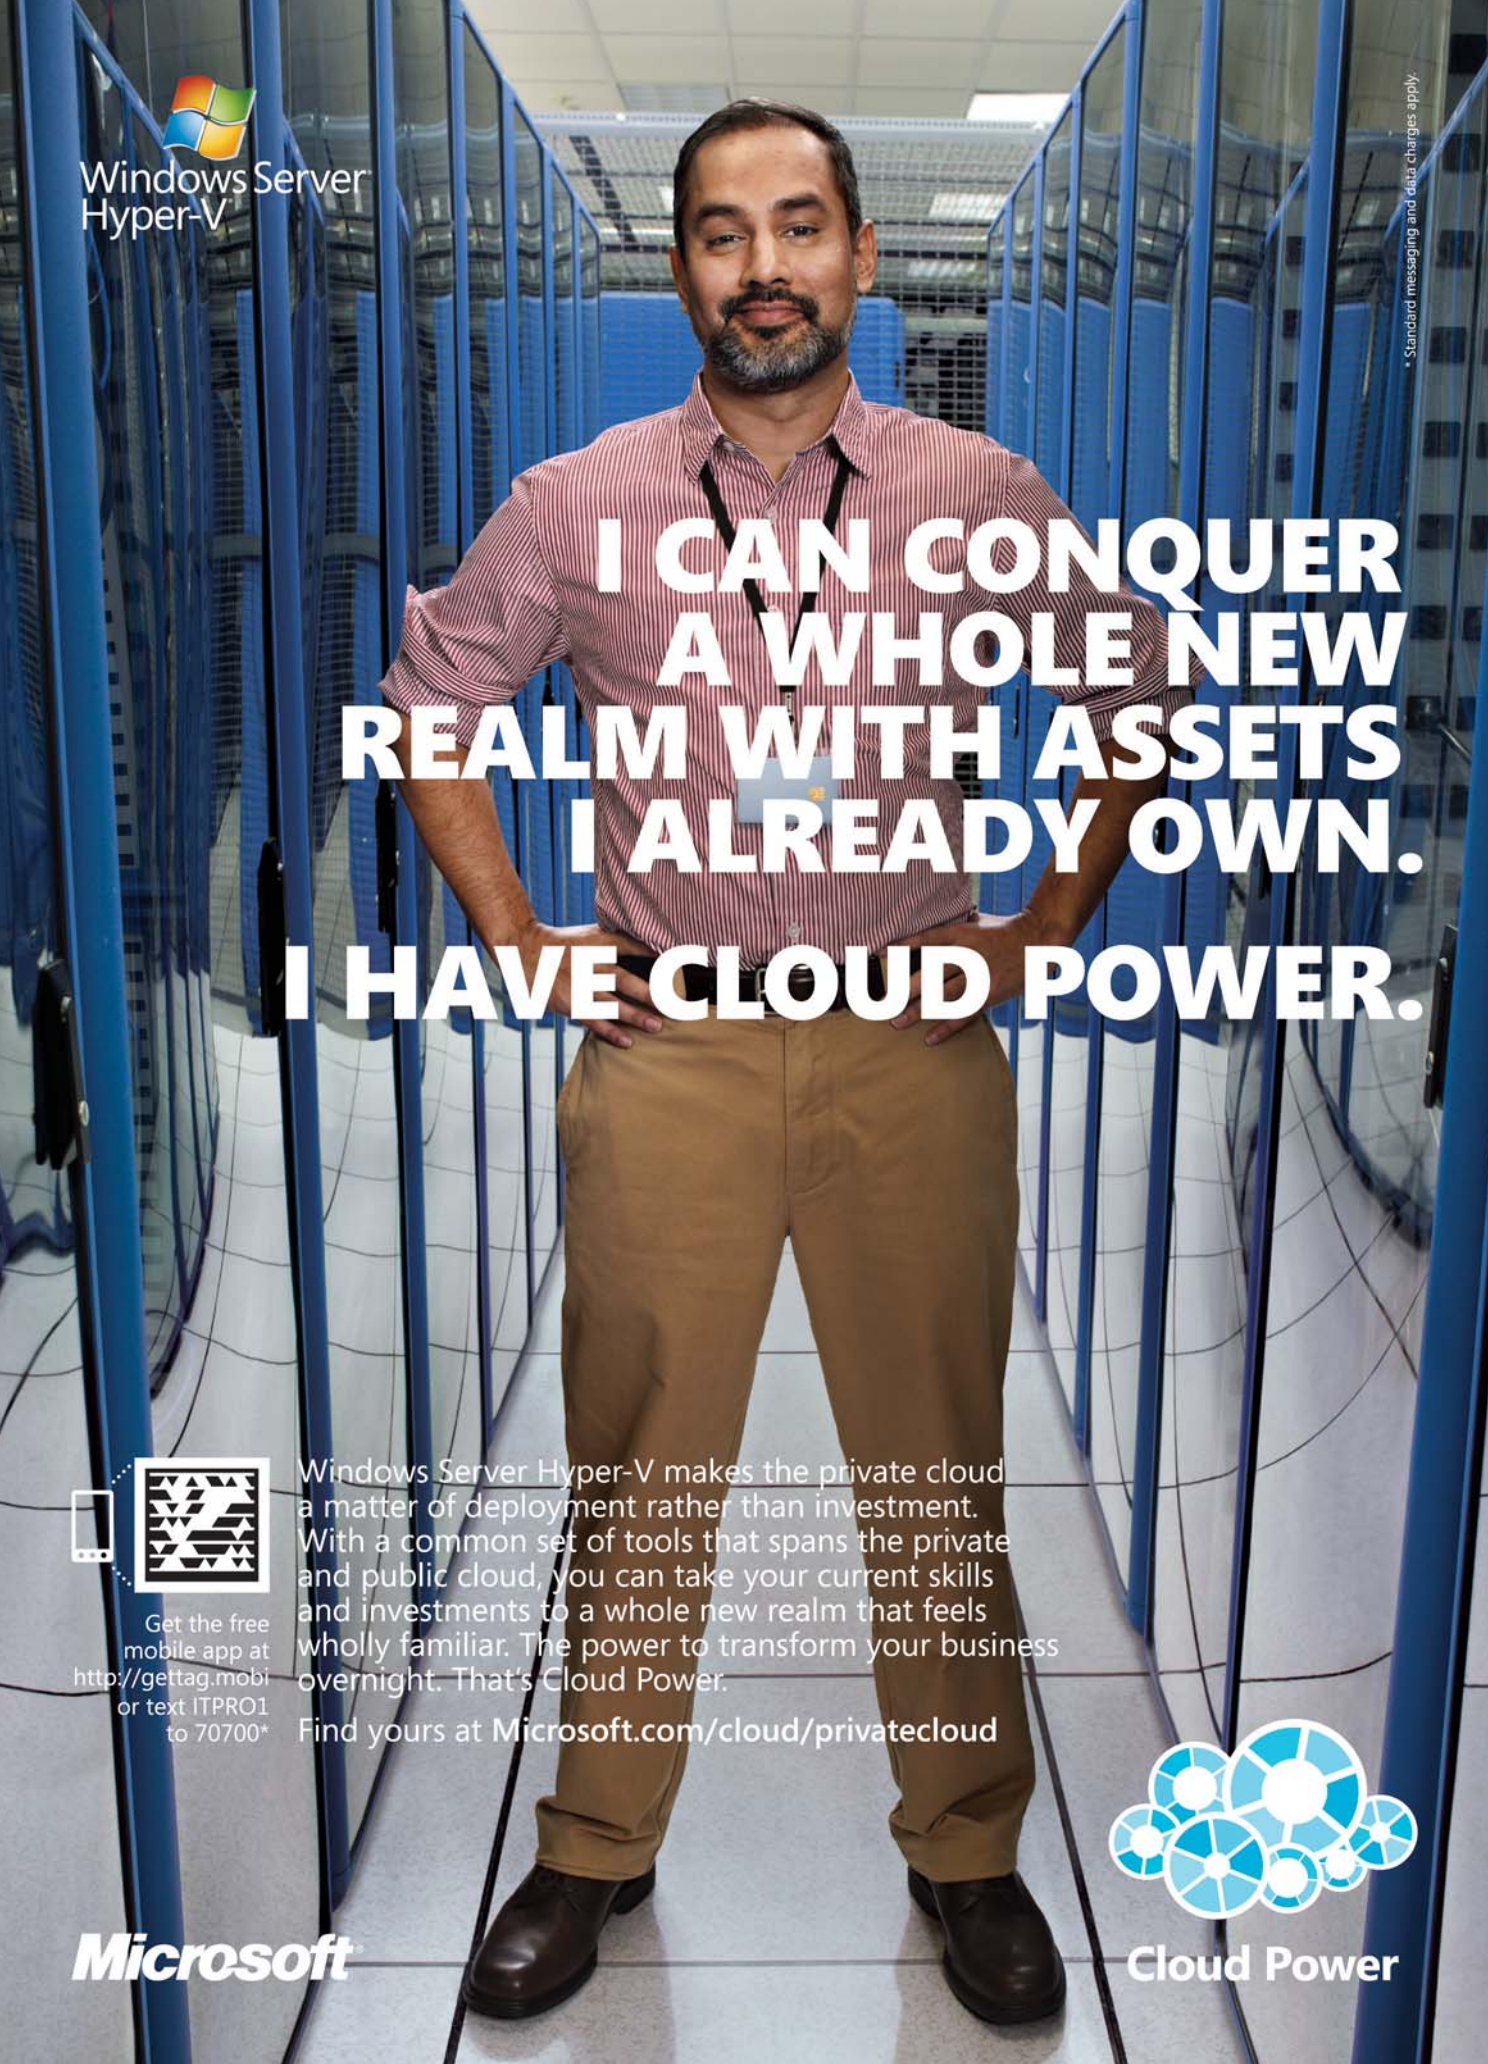
Windows Server Hyper-V
I CAN CONQUER A WHOLE NEW REALM WITH ASSETS I ALREADY OWN.
I HAVE CLOUD POWER.
* Standard messaging and data charges apply.
Get the free mobile app at http://gettag.mobi or text ITPRO1 to 70700*
Windows Server Hyper-V makes the private cloud a matter of deployment rather than investment. With a common set of tools that spans the private and public cloud, you can take your current skills and investments to a whole new realm that feels wholly familiar. The power to transform your business overnight. That's Cloud Power.
Find yours at Microsoft.com/cloud/privatecloud
Microsoft
Cloud Power